

**This Page Is Inserted by IFW Operations
and is not a part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images,
Please do not report the images to the
Image Problem Mailbox.

THIS PAGE BLANK (USPTO)



GB 99/01465

The
Patent
Office

PCT/GB 99/01765

INVESTOR IN PEOPLE

09/674717

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

REC'D 07 JUL 1999

WIPO PCT

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated

28 JUN 1999

**PRIORITY
DOCUMENT**

SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)


THIS PAGE BLANK (USPTO)

The
Patent
Office

Patents Act 1977
(Rule 16)

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)



9
05 JUN 98 E366162-1 003052
F01/7700 25.00 - 9812161.7
The Patent Office
Cardiff Road
Newport
Gwent NP9 1RH

1.	Your reference	A25547	
2.	Patent application number <i>(The Patent Office will fill in this part)</i>	05 JUN 1998	9812161.9
3.	Full name, address and postcode of the or of each applicant <i>(underline all surnames)</i>	BRITISH TELECOMMUNICATIONS public limited company 81 NEWGATE STREET LONDON, EC1A 7AJ, England Registered in England: 1800000	
	Patents ADP number <i>(if you know it)</i>	1867002	
	If the applicant is a corporate body, give the country/state of its incorporation	UNITED KINGDOM	
4.	Title of the invention	COMMUNICATIONS NETWORK	
5.	Name of your agent <i>(if you have one)</i>	Timothy Guy Edwin <u>LIDBETTER</u>	
	"Address for Service" in the United Kingdom to which all correspondence should be sent <i>(including the postcode)</i>	BT GROUP LEGAL SERVICES INTELLECTUAL PROPERTY DEPARTMENT HOLBORN CENTRE 120 HOLBORN LONDON, EC1N 2TE	
	Patents ADP number <i>(if you know it)</i>	1867001	
6.	If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and <i>(if you know it)</i> the or each application number	Country	Priority application number <i>(if you know it)</i>
			Date of filing <i>(day / month / year)</i>
7.	If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application	Number of earlier application	Date of filing <i>(day/month/year)</i>
8.	Is a statement of inventorship and of right to grant of a patent required in support of this request? <i>(Answer 'Yes' if:</i> a) <i>any applicant named in part 3 is not an inventor, or</i> b) <i>there is an inventor who is not named as an applicant, or</i> c) <i>any named applicant is a corporate body.</i> See note (d))	YES	

Patents Form 1/77

Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document

Continuation sheets of this form

Description	56
Claim(s)	5
Abstract	—
Drawing(s)	1

10. If you are also filing any of the following, state how many against each item

Priority Documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (*Patents Form 7/77*)

Request for preliminary examination and search (*Patents Form 9/77*)

Request for substantive examination (*Patents Form 10/77*)

Any other documents
(please specify)

11. I/We request the grant of a patent on the basis of this application.
Signature(s) Date:



6 JUNE 1998

Timothy Guy Edwin LIDBETTER, Authorised Signatory

12. Name and daytime telephone number of person to contact in the United Kingdom
Rohini R Ranjithkumar 0171 492 8146

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

Communications Network

The present invention relates to a communications network, and in particular to charging mechanisms in such a network.

5 In conventional communications networks, such as national PSTNs (public switched telephone networks), a significant proportion of the network resources are devoted to metering and billing network usage. Studies have estimated these resources as consuming as much as 6% of the operational costs of a telecommunications company. The Internet, by contrast, does not in general
10 incorporate metering and billing mechanisms for individual customers. The absence of the network infrastructure required to support metering and billing reduces the operational costs of the Internet compared to conventional telephony networks, and has facilitated the rapid expansion of the Internet. However the absence of appropriate billing mechanisms has significant disadvantages in terms
15 of the characteristics of the traffic carried by the internet: it encourages profligate use of network resources, and diminishes the incentive for investment in network infrastructure to support new applications requiring, e.g., guaranteed quality of service (QoS).

According to a first aspect of the present invention, there is provided a
20 method of operating a communications network including

distributing a tariff via a communications network to a multiplicity of customer terminals connected to the communications network, and

calculating, using the said tariff, a charge for use by the customer terminal of the network to which the tariff applies.

25 Reference to a terminal "connected to the network" is used here in the ~~description and the claims to encompasses terminals, such as mobile wireless data~~
terminals, which log on to a network temporarily, and other terminals which have a wireless connection to the network, as well as terminals which are permanently connected to a network by a fixed line. For example, a mobile terminal may log on
30 to a network to receive the tariff and subsequently calculate the charge while off-line, and such an arrangement falls within the scope of this aspect of the invention.

According to a further aspect of the present invention, there is provided a method of operating a communications network including;

distributing a tariff via the communications network to a multiplicity of customer terminals connected to the communications network,

at a customer terminal measuring use by the customer terminal of network resources; and

- 5 calculating, using the results of the said step of measuring together with the said tariff, a charge for use by the customer terminal of the network to which the tariff applies.

These aspects of the invention provide a lightweight charging mechanism suitable for use, for example, in the Internet, or as an alternative to conventional
10 billing mechanisms in other networks where the terminals have some data processing capabilities. It removes the burden of metering and billing from the network infrastructure and instead distributes the tariff to the customer terminals, allowing charges to be calculated at the edge of the network. This approach offers far superior scalability by comparison with conventional approaches, and is
15 therefore particularly suitable for use in a rapidly growing network such as the Internet.

Preferably the tariff algorithm is distributed to the multiplicity of customer terminals via the communications network to which the said tariff applies. In preferred implementations, the charging mechanism is designed to function even if
20 some tariff messages distributed via the network are delayed or lost. Preferably the step of distributing the tariff includes steps of communicating separately a formula for calculation of network usage charges, and coefficients for use in the said formula.

The network overhead for charging is further reduced by providing users
25 with the tariff algorithm and then updating only the relevant coefficients when the tariff changes.

Preferably the method includes measuring loading of network resources and determining a revised tariff in dependence upon the results of the said step of measuring loading.

- 30 A further significant advantage of the present invention is that it facilitates control of the use of network resources by amending the tariff to reflect the scarcity of a particular resource.

The steps of measuring loading and determining a revised tariff may be carried out automatically by a network management platform. Alternatively and

preferably, an algorithm for mapping congestion to price rises is distributed in the network, and preferably is located at customer terminals. Preferably the method includes operating a plurality of different services on the communications network, communicating different tariffs for different respective services to the multiplicity
 5 of customer terminals, and selectively varying a respective tariff depending on an operational condition of the respective service.

The different services may be distinguished only by different levels of QoS, or may be different in kind. This aspect of the invention may also be used in otherwise conventional networks, for example where billing is carried out centrally
 10 and tariffs are communicated to the end user only for information.

According to a further aspect of the present invention, there is provided a method of operating a communications network comprising:

- operating a plurality of different services on the network;
- communicating tariffs for the different services to a multiplicity of
 15 customer terminals via a common tariff distribution mechanism;
- and selectively varying a respective tariff depending on an operational condition of a respective service.

According to a further aspect of the present invention, there is provided a method of operating a communications network, including
 20 calculating for each of a multiplicity of customers, using a selected one of a plurality of different tariffs, charges for the use of network resources by a respective customer terminal attached to the network,

- measuring the loading of network resources, and
- varying one or more of the plurality of different tariffs in dependence upon
 25 the loading of the network resources, and in which different ones of the plurality of different tariffs have different respective volatilities.

This aspect provides customers with varying tariffs with different degrees of volatility. Then a customer needing greater stability can pay a premium to achieve that stability, while there still remains a band of higher volatility enabling
 30 the network operator to manage short term fluctuations in demand until longer term changes in tariff can be made.

According to a further aspect of the present invention, there is provided a method of operating a communications network in which at a point of access to the network a single blocking test only is applied to traffic entering the network .

Hitherto, a network such as the Internet has operated as a single service network. However it is now proposed that the Internet should become a multi-service network. For example, it may support multiple QoS levels for different applications, or might provide both multicast and unicast services to some but not all customers. The present inventors have recognised that, using conventional access control methods, this leads to a build up of multiple tests on access to the multi-service network to determine which service is being requested in each packet and then to check if it is a service which has been paid for by the relevant customer. This aspect of the invention overcomes this disadvantages by making a single blocking test that checks whether the customer is in a position to be punished for misuse of the network. Provided that this is the case, then the relevant packet is passed onto the network and all other appropriate checks are done in parallel, rather than blocking the packet while waiting for all the tests to be passed. If any subsequent tests are failed, for example if the packet has used a QoS level not paid for by the customer, then an appropriate punishment is imposed, for example by debiting a fine from a deposit lodged by the customer.

According to a further aspect of the present invention, there is provided a method of operating a communications network comprising:

- a) communicating tariff data to a user terminal connected to the network;
- b) calculating at the user terminal using the tariff data a charge for traffic communicated between the network and the terminal and making a payment;
- c) sampling part only of the traffic communicated between users and the network and for the sampled traffic comparing any payments made by users and the payment due according to the tariff .

According to a further aspect of the present invention, there is provided a method of operating a communications network comprising

- a) at a customer terminal measuring network usage;
- b) communicating network usage data from the customer terminal to the network operator; and
- c) the network operator sampling part only of the traffic communicated between a customer terminal and the network and for the sampled traffic comparing the network usage with the network usage data from the customer terminal and thereby detecting any discrepancy.

This aspect of the invention may advantageously be used in conjunction with one or more of the preceding aspects, but may also be used independently of them. For example, the customer terminal may measure usage data, and may send this data to the network operator, without having access to the current tariff.

- 5 The network operator might then apply the relevant tariff and bill the user based on the user's own data. In order to be assured that the network usage data is trustworthy, the data can be compared with the expected usage data based on the network operator's own measurements in a sampled period. If the data are identical, then the data for other periods is assumed to be trustworthy.
- 10 Alternatively, the tariff may be provided to the customer terminals and then, rather than the usage data being communicated explicitly, the customer calculates the usage charge. The payment of the usage charge, or equivalent accounting information is then communicated to the network operator, and the measured usage data is implicitly present in this communication.

- 15 According to another aspect of the invention, there is provided a method of operating a communications network, including automatically varying, depending on network loading as detected at a customer terminal, a tariff for network usage by a customer terminal. This aspect may be used in conjunction with, or independently of the other aspects of the invention.

- 20 Other aspects of the invention are as described and claimed below. The invention also encompasses communication networks, management platforms, routers and customer terminals adapted to operate in accordance with the methods of the invention, and computer-readable storage media bearing programs for implementing the invention in one or more of its different aspects.

- 25 Systems embodying the present invention will now be described in further detail, by way of example only, with reference to the accompanying drawings, in which:
-

Figure 1 is a schematic showing a network embodying the invention.

- As shown in Figure 1, a communications network 1 includes a number of
- 30 network sub-domains 2A-C. The network sub-domains may be under the control of different operators who may not trust each other. The network subdomains are interconnected by gateway routers 3, 4. In the present example the communications network is the Internet and supports both unicast and multicast Internet Protocol (IP) and associated protocols. A customer terminal 5 is

connected via a public switched telephony network (PSTN) 6 and an access router 7 to a subdomain 2A. A single blocking test is applied to traffic at this point of access, as further described in the attached paper. The gateway routers 3,4, and access router 7 may be commercially available devices such as CISCO series 7500 5 routers and CISCO series AS5800 universal access server respectively. Other customer terminals are connected to the network, including a Java-enabled mobile terminal 8 and a data server 9.

A network management platform 10 is connected to each subdomain. Each network management platform may comprise, for example, a computing 10 system comprising a SPARC workstation running UNIX (Solaris) together with network management applications. The network management platform 10 hosts management entities and tariff entities. The network management platform communicates with agents 100 in managed devices connected to the respective subdomain, for example using SNMP (simple network management protocol). The 15 management platforms monitor the loading of network resources in the respective subdomains, and, as will be further described below, adjust the tariffs for network use accordingly. The Net management platform (NMP) instructs the agent to monitor the device and report aggregated results at regular intervals back to the NMP, so the NMP can monitor the combination of all reports.

20 Tariff data is communicated to peer tariff entities in other subdomains and also to the customer terminals. The tariff data is multicast using, for example Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) dense mode. The tariff data channels are announced and monitored using protocols based on SDP (Session Description Protocol), SAP (Session 25 Announcement Protocol) Charging is carried out on a "pay and display" model. Each customer terminal monitors its own network usage, for example by counting the number of packets it sends or receives across the network interface. It calculates, using a tariff received via the network, the payment due to the network operator, and makes a corresponding payment into an account at the network 30 operator. The network operator polices the use made by customers of the terminal by intermittently sampling traffic to or from a particular customer and comparing the use made and the use paid for.

Preferred implementations of the invention are described in further detail, by way of example only, in the following technical paper.

Lightweight, End to End Usage-based Internet Charging: Architecture

[*Part I* | *Part II*]

Abstract

This paper describes a complete architecture for usage-based charging for any aspect of Internet service provision. The architecture is not another quality of service (QoS) mechanism designed for charging, but an architecture that can contain mechanisms for pricing and charging for QoS mechanisms and for other aspects of Internet service. It not only caters for end to end paths (both unicast and multicast) through the Internet across multiple network providers, but also caters for how the pricing and charging for these paths would combine in a whole network, even if some network providers do not operate usage-based charging. As well as dealing with network-wide charging, the architecture addresses the interfaces that link layer charging will present to network charging and the interfaces network-wide charging will have to present to higher level systems.

The architecture is intended to last for at least a decade from implementation (to claim longer is tempting but would be foolhardy). It is expected to exist in an extremely dynamic environment where the frequency of price changes for various aspects of the network service might only be limited by the spare bandwidth available to notify customers. On the other hand, the architecture can be optimised for the far more stable environment that is likely at the start of its life. The paper discusses how traffic volume growth is likely to be an order of magnitude faster than measurement system performance growth. With this in mind, the architecture is designed to survive the day when traditional traffic measurement even at the edge of networks becomes uneconomic. However, again, the architecture may be optimised for traffic volumes likely in the immediate future. In fact, a general principle for any particular aspect of the architecture is to identify the model which is the superset of all other models, whether it be pricing frequency, pricing granularity, payment frequency, who pays, how they pay, when they pay, who trusts whom, who needs accounting information, how often they need it, what is being measured and so on. However, generality does not imply abstraction. These general models are translated into real mechanisms that are simply chosen because they cater for all the requirements of more specific mechanisms. For example, if one buys 'planes with seats on runners, one can sell knee space by the millimeter, but one can also present a brand image built around three very spacious classes without changing the flying stock. About half a dozen of these mechanisms or models are believed to be innovative, mostly so that generality didn't lead to performance compromises. The paper is hence open to criticism and refinement.

The paper also contains analysis that is believed to be new, concerning the role of dynamic pricing in managing networks. It is proposed that there is no such thing as an elastic application, and that there is unlikely to be a need for admission control mechanisms in QoS protocols. The paper explores the role of dynamic pricing as a common denominator between federated systems to minimise management complexity. Pricing is also considered to be the best way to manage networks in the context of both the higher-level systems of which networks are a part, and the lower level components that networks consist of.

Keywords

Charging, dynamic pricing, tariffing, usage-based, flat-rate, process cost, micropayment, accounting, measurement, differential service, quality of service, fraud policing, protocol policing, punishment, admission control, policy, contract, pay-as-you-go, prepayment, spot pricing, congestion avoidance, demand management, class-based congestion, honesty box, pay-and-display, traffic warden,



repudiation, inelastic, time, operational support, dispersal, state, business component, business model, commerce, management, economics, separation, independence, architecture, open, end to end, end-end, scalable, lightweight, flexible, packet-based, connectionless, federated, Internet.

Contents

1. Introduction & Requirements
2. Justification
3. Architecture
4. Related work

7. References

1. Introduction & Requirements

This paper presents an architecture for charging for differentiated Internet service on a usage basis. The architecture is optimised for a connectionless network under highly federated management with frequent innovation in transmission techniques (viz. the Internet). Once it is accepted that some form of usage-based charging is necessary (justified later), a solution is required that neatly avoids (or at least addresses) the perennial compromises necessary between the following pairs of system characteristics:

- performance versus service flexibility (flexibility usually requires indirection layers which reduce performance)
- manageability versus service flexibility
- security versus accessibility (and possibly versus anonymity)
- security versus performance
- openness versus service differentiation

This introduction presents such a solution in overview, focussing first on performance by examining a lightweight end to end approach, then focussing on flexible business models to cover the other compromises, particularly analysing exactly what it is that a network provider might be charging for at the Internet service level. Having informally outlined the proposed architecture, the approach is justified in comparison with related work, starting by justifying usage-based charging itself, which includes a critical re-analysis of some of the accepted work in this field.

The body of the paper (separated into a second part) then rescans the architecture of this solution in more detail. This is broken down into sub-systems, which are described individually. Included in this is an analysis of how it might be possible to accommodate some of the major proposals of others within the architecture rather than just concentrating on the proposed approach as a monolith.

This leads in to a description of the related work in its own terms, also covering work on peripheral sub-systems as well as core charging proposals. This is followed by a brief exploration of the implications of taking the proposed approach as a precursor to identifying possible limitations, which are then analysed in order to build a programme of necessary further work. Finally, recommendations mentioned throughout the text are collected together and conclusions are drawn as to whether the proposed approach has indeed managed to avoid the above compromises or whether it is at least positioned correctly with respect to them.

The proposed solution is positioned to be in place some time during 2000 to 2002 and to have the flexibility to evolve and to support the evolution of higher level systems for at least a decade.

1.1 Lightweight, end to end approach

This paper investigates a shift in approach to usage-based Internet charging with decisions centred around end-systems not the network. The primary motivation for this paper is that the introduction of charging itself brings its own costs; usage-based charging particularly so. Varian quotes billing and accounting as about 6% [Bailey95] of total costs of an unnamed telecommunications company. These costs will have to be borne by all users of the technology.

ne term "charging" is used throughout to describe the whole process of offering a service for money and accounting for its use and settling. This is distinct from some definitions, which use the term in the narrower sense of chargeable event collection, independent of pricing. Rather than describing the problem further, then describing the approach to solving it then describing the solution, we will dive straight into a broad-gloss description of our solution and justify it later.

In the proposed approach, the network provider no longer does billing and accounting for its customers, only monitoring, pricing and policing, because the customers do accounting for themselves and for their provider. This applies to both edge and backbone network providers all of whom base their accounts on measurements taken at the end-systems that are aggregated by the accounting processes of each network provider as we move deeper into the network. Customers contract with their next level network providers to measure their own use of the network (or have it measured for them) and pay into "honesty boxes" accordingly. The network providers contract to provide tariffing to support the users in this. Network providers may all operate different tariffing regimes, some even providing services for free that others charge for. This is possible because the approach involves no changes to transmission protocols, instead being implemented in parallel to the transmission system, only connected through references in dynamic contracts. The new information infrastructure needed to allow these contracts and tariffs to be implemented efficiently is described in architectural terms. The advantage in responsiveness for managing demand compared to network-based billing is explained.

In this move to a "pay-and-display" model, the network providers protect their interest by operating "traffic wardens" that sample their customers' traffic, gathering evidence of possible fraud. The proportion of each customer's traffic that a warden monitors may be varied depending on the level of trust in that customer. In the extreme, 100% of a customer's traffic may be monitored which emulates a more traditional billing approach. Thus this scheme is a superset architecture of traditional billing (this scheme is designed to be a superset of other approaches in many other aspects too, which are described under Flexible business models). A form of contract is proposed which would allow network providers to de-risk this situation by pre-agreeing the punishment that the customer accepts may be exacted given specific evidence of fraud. This is important to reduce the cost of debt recovery and complaint handling. A typical example might be a deposit which the network provider has a right to draw on in specific circumstances, the balance of this deposit being connected to the user's ability to access the network for some or all classes of traffic. From the customer's point of view, as long as they are using network access software that is accredited by their network provider, the expectation is that they will never fall under suspicion of fraud.

This system of monitoring and punishment is independent of the tariffing and accounting systems and may therefore be used to police incorrect use of standard network protocols even where no usage-based charging is involved (e.g. to punish contravention of aggressive use of TCP). The contract between customer and network provider would obviously have to stipulate compliance with such standards (possibly referring to some software accreditation scheme) and highlight appropriate punishments.

Figs 1a-c represent the main elements of the system in three phases:

- a) set-up
- b) operation
- c) monitoring (during operation)

The hexagons represent the administrative domains of the Internet managed by different parties. The edge-customer of the edge network provider is represented by the collection of clip-art in the bottom left corner. The intention is to reinforce the general nature of the term "customer", which may be an individual or a corporate body and may be using Internet-enabled devices that are either general-purpose computers or embedded systems for specific applications. The "server" clip-art in Fig 1b highlights the fact that the payment system may be separate from the system using the service - not only just physically separate but also possibly spending another party's money. The piggy-banks represent the financial systems of the network providers. The thin straight arrow represents utilisation of the Internet by end-systems sending or receiving packets of certain classes in certain patterns. Accounting between network providers involves an extension of the same "pay and display" model, but will be covered later in this introduction.




Fig 1a - Pay and display model: set-up phases

The first part of the set-up phase (Fig 1a) is when terms and conditions (Ts & Cs) are agreed resulting in some redress for the network provider if the customer breaks the terms. This redress may be denial of future service, blacklisting of the customer, or, as in the case shown, a fine of an agreed amount taken from a deposit (to reduce the cost of fine collection). Incidentally, where a deposit system is used, this in itself helps the funding of investment in Internet infrastructure. Indeed, the deposit could be considered as a "share in the Internet" (whether capitalist, common ownership etc.), given that providers can bank on the fact that they won't need to draw on all deposits at once.

For the second part of the set-up phase, the Ts & Cs would also constrain the customer to agree to run software that would monitor and act on electronic announcements on current tariffing etc. and the provider would agree to maintain this capability. Tariffing information would be updated in different epochs requiring different mechanisms (described later) for different announcement frequencies and reliabilities.

These tariff channels are represented by the curved solid arrows in Fig 1a. The main information that would be disseminated in these announcements (in roughly increasing order of frequency) would be:

- contract variations
- new payment interfaces or removal of old ones
- new chargeable services with their tariff structures or removal of old ones
- new instructions on which accounting information the provider requires from the customer for each service (not shown)
- new code to measure these new services or recommendations to remove old code
- new tariffs for existing services (e.g. conditional discounts, cross-service dependencies)
- algorithms defining the mapping between measured congestion and the price ("fine") to be paid for reacting to it abnormally
- regular price updates within existing tariff structures that could be infrequent or highly volatile (spot pricing)
- possibly, regular updates of the dependence between price and remote address

The curved dotted arrows in Fig 1a represent the references to the payment interfaces of the network providers' financial systems. It should be noted that it is possible for a number of network providers to share one payment interface or one provider to announce many payment interfaces.

Note that it is possible for the customer to be using chargeable services after installing measurement software but before listening to tariffing information. This is an important point, in that it enables low latency, serendipitous access to any chargeable service (for which software is installed) even if the

12
 Trifling information takes long to be accessed. This is on the assumption that the customer is willing to risk using a service where the current price is unknown (a previously cached price may be sufficient to reduce this risk).

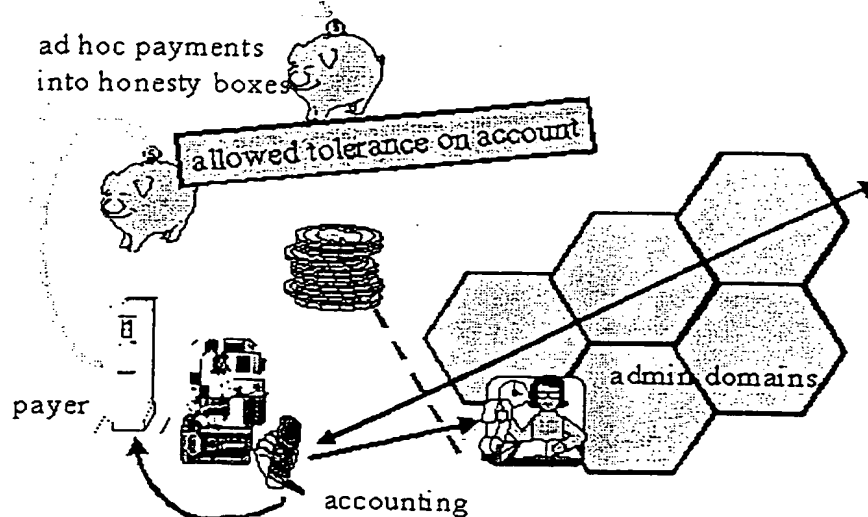


Fig 1b - Pay and display model: lightweight during operation

Fig 1b shows the operational phase when the end-user is using the network and paying for this use as two decoupled operations on two independent systems. Each customer will have an allowed tolerance on their account taking into account their deposit and their credit-worthiness that will probably mean that payments can be fairly infrequent. However, as long as a low-cost payment medium is employed (micropayments), the architecture allows for very frequent payments if necessary (e.g. where another party of unknown trustworthiness is paying on the customer's behalf). The device using some chargeable class of Internet service is sending messages to the relevant payment system instructing it to pay as required. The edge-network provider's contract *may* also require that the end-system sends accounting information regularly to the edge-network provider. Again, it may be stipulated that this accounting information must be supplied as service is consumed, or more typically aggregation into less frequent messages might be allowed.

As promised earlier, the relationship between neighbouring network providers needs consideration. In general, pay-and-display and random traffic warden policing can be used recursively to calculate the wholesale charges to be transferred from edge network providers towards the backbone providers or towards directly connected peer providers. An edge provider can build an aggregate database of all the accounting detail (including source and destination address) supplied by its customers for a slice of time. Taking a time averaged view of its border routing over this period, these accounts can then be allocated to each of the links with other network providers. The sliced up and aggregated accounting database can then be passed over to the relevant neighbours so they in turn can continue the process. In fact, pay-and-display is the standard model [ITU D.150] for accounting between international carriers in the public switched telephone network (PSTN). Experience from this and more detail on inter-provider accounting is given later.

The payer may not be the end-user but some other party (e.g. a video on demand service which bundles network quality of service charges). In such cases, the payment system may require proof that the service has been consumed within the terms it has agreed to pay for, which will involve the end-system requesting receipts in return for its accounting information which it can present to the payment system. These may be required all the time, or only when challenged randomly. This scenario is dealt with in more depth later.

It can be seen that, during normal operation, the charging system is typically very lightweight being separated from the system consuming network service in terms of time, space and identity with no

customer-specific measurement taking place in the network. However, this doesn't preclude bulk statistical measurement in the network to determine pricing variations. The tariffing channels are still assumed to be feeding price variations to all end-systems during operation, and the end-systems must be listening under the terms of their contract.




Fig 1c - Pay and display model: cheats get punished

Fig 1c shows the situation when, during the operation phase, the edge network provider decides to enter a policing phase for a particular customer. A "traffic warden" process will gather usage statistics from the network interface with the customer and compare them with the customer's payment accounts over a sufficiently long period to satisfy itself that the customer is keeping to the contract. This is not sampling in the sense normally used in network measurement circles [Bohn93, Ruth97], where the sample measurements are taken to be representative of unmeasured periods and factored up accordingly. The idea is simply to check that end-system measurement is correct during the sampling period then trust it during periods not sampled. The sampled measurements themselves are discarded as they are of little use in predicting unsampled usage due to the bursty nature of data traffic.

If it is discovered that payments are falling behind usage in a systematic way, the contract will specify what evidence is necessary to be collected before the right to punish the customer is assured. The network provider may also operate business rules not in the contract (e.g. to be lenient to major accounts or infrequent offenders, to give people three chances etc.). In the diagram, the case is shown where the agreed punishment is to dock some of the deposit. In this case, the customer's network access is dependent on the deposit level. Typically, this would merely cut access to use-based chargeable services, or the service(s) not being paid for, leaving some base "best-effort" flat-fee service intact. The rationale for this might be that the customer is still up to date with their subscription for basic network use (and it wouldn't cut off the customer's ability to settle their debt by an on-line payment method!). Note, though, that the network normally deals with network addresses not the identity of individuals using these addresses. The security of the dynamic mapping between these is discussed later.

It should be noted that barring certain traffic at the customer interface would have to cover both incoming and outgoing packets. This would still leave the network provider using resources for packets destined for receipt by that customer rather than sent by the customer. As these cases would be exceptional, this is unlikely to be a problem, but otherwise, this points to the need for send barring protocols in the Internet, which would fix many of the vulnerabilities to denial of service attacks, but would also be open to abuse if not designed carefully. An alternative is explored later.

Obviously, the level of punishment and the frequency of sampling would have to provide sufficient deterrent to fraudulent users. Also the evidence gathering process itself would have to be undetectable and random, otherwise users would simply start behaving themselves when they knew they were

and random, otherwise users would simply start behaving themselves when they knew they were being watched. This is discussed further later.

Before moving on, it will serve to emphasise an important architectural principle that has been quietly introduced. During normal operation, each packet need at most be subjected to one blocking test (Fig 2). That is, the edge network provider need only block a packet while checking that the customer has lodged some token of trust (e.g. the deposit or an account contract). All other tests and policing may be done as parallel threads to the forwarding of the packet. Thus, one can avoid checking in some serial sequence of tests that the customer is entitled to use this service or they have paid for that service, or that, when they say someone else will pay, the contract they have with that other person can be checked. Instead only one commercial test is necessary in order to allow all these other tests to proceed asynchronously while allowing the packet through. The reason that this is important is that, currently, in a world without differential IP services, access control is applied on entry to the whole IP service layer and is generally provided and revoked on long time-scales. If access is to be granted or denied to individual classes of IP service, this could imply that every packet is to be checked against a list of rules at the access gateway, hence the need for this principle. This should avoid the situation in the similar scenario of an IP firewall, where every rule in the access control list can reduce throughput by about 10% [ACL perf?]. Alternatively, if there is only one primary customer per link, a far more efficient arrangement is to only do a blocking test if the token of trust becomes insufficient.




Fig 2 - Single blocking test

To clarify, the single blocking test is very different from Clark's in/out bit [Clark95], which indicates whether a packet is within the terms of the contract with the customer or not. No assumptions are made here that the customer needs to agree a contract over their usage profile of network services. This is merely a test as to whether there is some incentive or deterrent present that will enable the network provider to enforce any contract, whether to do with usage profile or other contractual issues like agreeing the meaning of tariff announcements. Consequently no bit is necessary in packets to show the test has been passed as it only concerns the commercial relationship between the two parties across the customer interface. Typically, there will be no need for any dynamic blocking tests between ISPs as their relationships are more long term and better dealt with out of band.

To summarise this introduction, the proposed architecture removes the vast bulk of the responsibility for event collection and billing from the network provider, placing it on the end-system, which it keeps up to date with dynamic contract channels. We now move on from the performance aspects of this proposal to the aspects that improve business flexibility *without* compromising the lightweight aspects.

1.2 Flexible business models

Flexible business models

1 What business?

Before discussing flexibility of business models, we must be clear what business we are talking about: operating as an Internet service provider (ISP), either as an edge network provider (EN) or as a backbone network provider (BN).

In the context of this paper, this means offering end-customers a connection to equipment which can understand the Internet protocol (IP) and the Internet control management protocol (ICMP) and arranging relationships with other ISPs so that any IP packets presented to us can be forwarded to the correct exit point from our network. These inter-ISP relationships consist of mutual connections, agreement to swap routing information and application of policy over routing (which may involve charging). In the case of IP multicast, an ISP also operates routers that duplicate packets and forward them to multiple exit points based on multicast routing information driven by Internet group management protocol (IGMP) [RFC2236] requests from hosts.

Offering global connectivity implies offering differentiated connection lengths. Invariably, the extra cost of longer routes is subsidised by over-charging for shorter routes to avoid the cost of differential charging. Currently, the extra latency packets endure over longer routes is considered sufficient incentive to encourage demand to tend towards shorter routes such that it approximates to supply without a charging mechanism [Caimcross97]. Fortunately (in this respect), the speed of light is likely to remain an insurmountable barrier to lower latency communications for the foreseeable future.

Not charging differentially for distance is an example of a more general principle: a network is offered to customers as a "black box" with defined interfaces, similar to an object in an object-oriented software system. Thus, if the internal implementation of the network is not optimal in all scenarios, the customer isn't charged more when they happen to make a request that proves difficult purely because of internal design decisions. If two connections to the network are physically close, but the shortest route between them within the network goes half-way around the world and back, the customer isn't penalised in price terms. Similarly, if the subscribers to a particular multicast session happen to be located in such a way that the multicast tree fans out completely at its entrance to the network, it is not the customer's problem that this happens to be no more efficient for the network provider than multiple unicasts. If a network provider decides it is more efficient to operate simple routing protocols than ensure every route is optimal, the network provider has to accept the consequences. This is not to say that pricing can't depend on remote address. Rather it shouldn't depend on the route between the two. This is an incarnation of edge pricing [Shenker96]. The distinction is that routers can make decisions over the best route to a destination based on cost metrics, but end-systems may only make decisions on which other end-system (where a choice exists) they would rather communicate with and when they would rather do it. This only works if end-systems can assume routers are acting in their best interest, which is only a safe assumption in a competitive or well-regulated environment. We would hope to avoid the clumsy interruptions or access codes that regulators have forced on end users in some telecommunications markets to give them a choice of trunk carrier.

~~Running an ISP business does not mean the ISP has to own or operate the links and link layer equipment over which some or all of its network is built. These can be run by separate businesses who will charge for their use on a link by link (and possibly usage) basis. Similarly for the cases when IP traffic is routed by non-IP network layer technology⁽¹⁾ (effectively tunnelled) although the commercial relationship here is more complex (see Further Work). These charges are only relevant to this paper in as much as they contribute to the cost base of the ISP business and therefore to the calculation of the charges levied for Internet service. Thus, costs of link layer or non-IP network layer infrastructure and any connection charges to other ISPs are simply an ISP's wholesale running costs. The value an ISP adds to its wholesale purchases is essentially incremental extra connectivity.~~

An ISP cannot sell a transport layer connection *per se* as a service, because it doesn't supply a connection - the end-systems create and manage the connection using TCP and their owners would be justifiably piqued if they were charged for something that they had to produce themselves. However, it will be described later how the management of a connection is effectively a contract between the end-systems and the network, and various schemes have been proposed for selling the right to vary this contract to the customer's advantage.

16

Although an architectural principle of the Internet is "do one thing and do it well" [Huitema95], connectedness and routing is not all an ISP can offer. An ISP business can increasingly operate equipment that also reacts to aspects of IP and its related protocols that demand differential Internet services. Differentiation can be in terms of latency, bandwidth or reliability (either absolutely or in terms of their derivatives with respect to time), collectively termed quality of service (QoS)⁽ⁱⁱ⁾. Further, this equipment must contribute to an Internet-wide (end to end) differentiation of service for such traffic through compliance with standards.

Here it should be noted that the responsiveness and reliability dimensions of QoS are provided additively by a chain of ISPs on an end to end path, while bandwidth is provided dependent on the minimum (bottleneck) ISP. For the major classes of traffic where usage-based charging is being applied for reliability, it is likely (i.e. recommended) that, unlike for best-effort or controlled loss traffic, packet drop will be sufficiently rare⁽ⁱⁱ⁾ that peering arrangements will explicitly agree to mutually ignore it. This is on the assumption that, where reliability is being paid for, ISPs will only interconnect where they both have sufficiently dimensioned networks to avoid dropping such packets in all but exceptional circumstances. A similar recommendation applies where latency rather than reliability is concerned - if responsiveness is being charged for, it would seem more efficient to assume all ISPs on the path are dimensioned sufficiently to provide the requested latency in all but exceptional circumstances. Further the ISPs could agree that no-one receives revenue in the unlikely event that the requirements sometimes fail to be attained. The alternative would be to develop means to prove who was to blame for dropping a packet or delaying it more than some "fair" share of the total latency budget - both very difficult problems to achieve between mutually untrusted parties. This "exception peering" would take peering agreements to a more subtle level, with a consequent need for background (sampled) monitoring of neighbours to check that they weren't gaining unfairly from such arrangements by insufficient dimensioning.

Differential service could be achieved by operating a number of networks and allowing customers the choice of which to use. These differentiated networks could be physically separate or more interestingly only logically separated such that a customer could request to switch to a different service level and be given or denied long term access to a different level in fairly short time-scales. Theoretically, one business could supply the basic Internet service to another who would add differentiated service⁽ⁱⁱⁱ⁾. However, there are considerable advantages to operating multiple levels of service on one network on a packet by packet basis. Customers are likely to prefer this rather than paying a premium for access to a high-class network when not all packets need it. Network operators are likely to prefer it as it involves managing one network [Crowcroft1n?] and allows the higher class provision to absorb bursts in demand by eating into the lower class provision [Floyd?]. A brief "proof" that two classes on one network is more efficient than separate networks is provided by Shenker [Shenker95].

Often ISPs operate related Internet services (domain name service, application layer caches etc.) and related communications services (multi-protocol routing, link layer services for themselves or others) but these are not the business that is the subject of this paper.

To summarise, for the purposes of this paper, an ISP provides:

-
- access connectivity
 - global connectivity through interconnect
 - packet forwarding based on routing and routing policy
 - access bandwidth up to the physical limit of the line equipment used
 - differentiated route length implicitly included in global connectivity
 - packet duplication based on multicast routing and policy
 - differentiated packet forwarding in terms of end to end latency, bandwidth and/or reliability

Access connectivity is sometimes charged on a per-time basis by the link provider or by the ISP, with some cases where both are charging and others where neither does^(iv). Global connectivity and forwarding over it are currently charged on a flat-fee, subscription basis (you either have connectivity or you don't), at a level related to the maximum access bandwidth of the connection to the ISP.

Usage-based charging for variable access bandwidth (up to its physical limit) would be a refinement of this current scheme. Per-time charging for connectivity could be subsumed into such a scheme (by

14

ing a positive price for (but connected) bandwidth). The techniques for usage-based charging discussed in this paper could be applied to charge for variable access bandwidth but their real strength is that they can also be applied to the differential services (QoS and even route length) and to multicast^(v). It then makes sense to use these techniques for variable access bandwidth rather than having to manage two charging systems.

Later, when the architecture is described in more detail, it will briefly be explained how various schemes proposed in the literature for chargeable QoS or for pricing could be implemented efficiently within this architecture.

As well as the primary services listed above, an ISP offers a degree of packet-level management control (using ICMP), but this is more likely to be used to support charging, than be charged for itself (although the proposed scheme could do this). An ISP may wish to provide some degree of network management control or monitoring to a customer. This can also be charged for on a usage basis with the proposed scheme.

On further analysis, it is possible to identify further second order services that an ISP could provide (and charge for), such as stability of end to end bandwidth provision, stability of price, smoothing of bursty traffic to a server etc. These are considered later.

Just because any of the above aspects *can* be charged for, doesn't mean it makes business sense to charge for them all. It does make business sense to have a charging system that can be applied to any business model at any time if this can be done efficiently for each specific case.

1.2.2 Sender or receiver liable?

Above we introduced the concept of an ISP's network being similar to an encapsulated object. Before moving on, while on the subject of "what business" an ISP is offering, a couple of common confusions need sorting out concerning the direction of transmission of data across this object's boundary.

The first confusion is to do with terminology. Certain QoS mechanisms (e.g. resource reservation set-up protocol (RSVP) [Zhang93]) involve the receiver of some data initiating signalling to set up a reservation for the QoS of subsequent data reception. Thus the "data receiver" is the "sender" as far as charging for the resources requested is concerned, because they are sending the signal that causes the costly reservation. This signal will typically cause resources to be consumed all along the end to end path (implying the "signalling sender's" ISP, the "signalling receiver's" ISP and backbone ISPs will all be looking for someone to call to account for the cost).

If, instead, certain flags in each packet determined the level of resources to set aside for them, the data receiver would be the "signalling receiver" because this is per-packet signalling. If some third party set up some mapping in the network between a level of QoS and certain flags that might appear in packets, then this third party would be the sender as far as charging for such set-up signals is concerned.

The second confusion concerns the difference between charging for the resources used by sent messages as against those for received messages (whether data or signalling). A large body of the literature is very sloppy in this respect, even when only considering fixed access charges. Many authors (who should know better) state that they believe the current Internet model is "sender takes all" [ITU96, Zull97]. That is, the only revenue received for access bandwidth is by the sender's ISP, as if traffic out of a network doesn't need any bandwidth. Very simply, from a cost point of view, there is typically no distinction between sent or received traffic to the ISP. Whether or not the sender or the receiver is responsible for setting some level of QoS on a packet, it still costs the ISP to transmit it across its network whichever direction it is travelling, including across the interface with its customer and across the remote interface(s). The problem of customers receiving unsolicited packets is not one to be solved by only charging the party that originally solicited, because such "blame" is invariably impossible to determine at the network level [MacKieVar92]. For unicast^(vi) traffic, this seems to point to a requirement (that is not easy to satisfy) for a "call-barring" protocol; that is, a way for all parties on the Internet to be able to request their upstream neighbour to bar certain traffic types from certain source addresses destined to their target address. However, a possible commercial model is described below that might avoid this need.

Were an ISP to assume that all incoming packets resulted in exactly one outgoing packet, it could make a sweeping generalisation and cover sending costs with reception charging (or *vice versa*). However symmetry is not the case (due to packet drop^(vii), multicast and aggregation^(viii)), so it is recommended that customers at every interface to a network should be charged for transmission in either direction. Considering packet drop first, charging only on receipt would ignore the cost of transmitting packets before dropping them, and fail to give the sender an incentive to adapt to packet drop. For multicast and aggregation, charging senders and receivers ensures costs are covered without having to introduce mechanisms for measuring the fan-out of a distribution tree (which was avoided in the design of such protocols to keep them scalable by constraining the effect of a join or leave to be local). To clarify, this doesn't preclude everyone's charges for multicast or aggregation ultimately being covered by a different sub-set of the parties (e.g. multicast senders only), but this would be through higher level transactions (see below). Although the point of this paper is to allow any business model, rather than recommend particular ones, this is a case where the alternatives are more complex or wouldn't be stable (anyway, the recommendation is only from a technical point of view and can be flouted if commercially sensible).

A consequence of such a rule *appears* to be that all hosts have to be ready for any standardised QoS mechanism to arrive in packets received and to account for them correctly. Further, it *seems* necessary for bodies like the Internet Engineering Task Force (IETF) not to standardise too many QoS mechanisms. However, as promised earlier, an ISP can adopt an even more subtle approach to charging for reception. It can make it *customary* for the receiver to be charged but put the ultimate liability on the sender. A similar but opposite situation used to prevail with the UK postal service. It was customary for the sender to pay for the stamp, but if it was missing or insufficient the receiver was liable for the payment whether or not they wanted the letter because the Royal Mail had an obligation to deliver every letter^(ix). Thus an ISP could create the contractual situation where a receiver is normally expected to pay, but if a one refuses it is incumbent on the sender to prove they were asked to send, or pay up. ISPs would have to include terms in their contracts with other ISPs making them liable for passing on unsolicited data rather than the ultimate sender, otherwise it would be impractical to take action against another ISP's customer who might be anywhere in the world. This would encourage ISPs to include terms controlling their end-customers, so that they could charge them for unsolicited sending or even kick them off if they persisted. ISPs already have similar arrangements for e-mail spamming which are successful to varying degrees^(x).

While on the subject of asymmetry, it should be noted that the prices for QoS (or any other aspect) assigned to transmissions into and out of a network may not be the same. For instance, with an asymmetric access technology like asynchronous digital subscriber line (ADSL) or satellite, it may well be more costly to provide QoS in one direction than another, and consequently it may be priced differentially. Another example is multicast where the QoS of sent traffic could be charged at a premium despite no difference in underlying costs, but simply because senders are ripe for exploitation until the reduced costs of multicast have worked their way through the market.

These recommendations make more sense in the context of how they would be applied for various known QoS mechanisms that are each either sender or receiver initiated. Thus, further discussion is left for later, as is discussion of how to treat caches and proxies.

To summarise, provision of network service always emanates outward from a network provider, irrespective of the direction that data or signalling is flowing. The consumer of these services is the directly connected network or end-system, who is considered liable for charges in the first instance. The exception being that, when the consumer claims resources were consumed due to another party sending unsolicited packets outside its control, the directly connected sender is ultimately liable in the absence of proof to the contrary. It must be made clear that these discussions only concern who is *primarily liable* for charges for resources used. Who actually settles this charge is another matter, dealt with next.

1.2.3 Internet business as a component

The aim must be to provide Internet service with usage-based charging in such a way that it can be used as a component of any higher-level information system. By "information system" we mean any service being offered on top of the network layer services as defined above. This might be an application layer service such as Web, databases, control systems, directories, conferencing, telephony, gaming etc. It could also mean services which are distributed but fulfil the role of layers of

19

phony, gaming etc. It could also mean services which are distributed but fulfil the role of layers of OSI stack above the network layer but below the application (e.g. specialised encoding hardware, encryption, local repair of multicast packet loss etc.). In order for ISPs to offer their customers (those building these higher level systems) a simple but comprehensive business interface onto usage-based charging, they need a simple abstraction of their customers' needs.

This boils down to three scenarios shown in Fig 3:

- applications where there is no charge for the information, only for the communication (e.g. conferences^(xi), telephony)
- applications where the information and communication are charged independently (which is the same as the first scenario from the network providers point of view)
- applications where the communications charges are bundled into the information charges

It would be possible though unlikely for an information service that charged on a flat-fee basis to incorporate usage-based transmission charging.

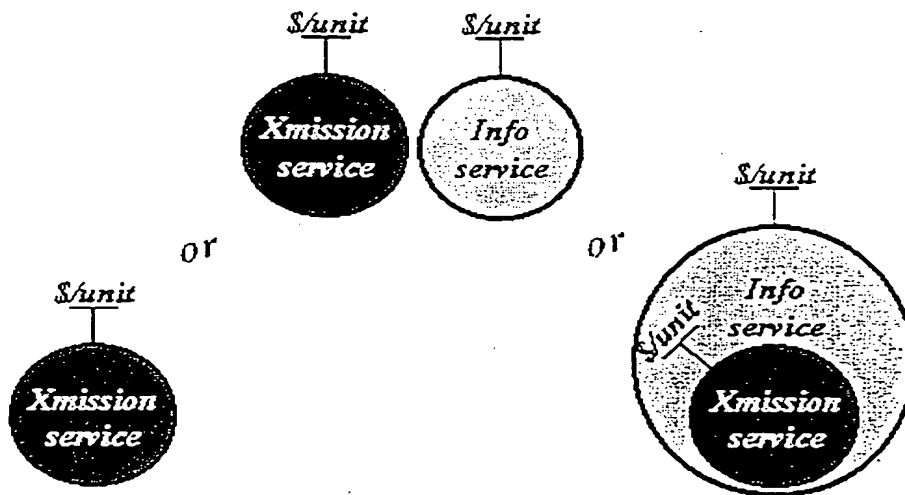


Fig 3 - Usage-based charging for transmission: separate or integrated with information charging

For instance, a video streaming service might operate under either of the last two models, but would require very different charging systems in each case. In particular, if the communications charges were bundled, it might be very important to the service provider to ensure the money it was paying for each customer's transmission quality was indeed achieving good transmission. It may even be necessary to stop a customer redirecting this money into improving the quality of some unrelated transmission to the detriment of their video service. In particular, it may be a condition of advertisers, that customers shouldn't be able to do this redirection of money for quality during their adverts.

So, in order for Internet service to be easily integrated as a component of wider services we end up with different charging system requirements dependent on whether:

- the consumer of network service is the same as the payer
- and if not, whether the payer needs the consumer to prove that the network service was used for transmission of specific information

Clearly, all ISPs support each other in providing global connectivity and will have to co-operate to provide various levels of differentiated global service. Therefore the business component for transmission service must itself be capable of encapsulating the businesses of all the ISPs on a (possibly multi-ended) path. Thus the need to offer network service as a simple component is no less important when considering the service one ISP supplies to another. However, each one will want to be able to offer its service as a component of its customer's higher level systems without that

- provision of a service, A. that "bundles" provision of service, B. supplied by a "sub-contractor" where the customer has no relationship with supplier B and so the price of the sub-contracted element can legitimately be different from the cost which can remain secret
- provision of a service, A. that "covers" provision of service, B, supplied by a "peer-supplier" where the customer has or could have a direct relationship with supplier B and expects service B at the same price they would pay directly to B

In order to cover all these scenarios, it is necessary to conform to some general principles:

-

It has already been stated that, just because the proposed scheme *can* charge for all sorts of aspects of an ISP service, it doesn't mean it *must*. It is also recognised that a system that is very flexible can sometimes be inefficient at specific tasks (Jack of all trades, master of none). It is therefore necessary to recognise that the proposal in this paper is initially being used to build a test-bed system where flexibility is paramount. The experience gained will feed into operational design, where it may be sensible to remove some of the flexibility. However, where possible, flexibility is included in such a way that it is only invoked during set-up phases, not during regular operation.

21
particular example of this where contracts are used to connect the concepts in the charging system to the concepts in the transmission system (Fig 4). This should, in most cases, allow Internet protocol definitions to remain unchanged as usage-based charging is introduced or changed (that is, assuming protocols handle QoS, multicast or whatever they don't need to be altered to charge for it).

It should be noted that the charging system would, to a great extent, be an on-line system and use the transmission system to transport its messages (including channels for the contracts themselves as described early in the introduction). However, at the network layer, these charging messages would be no different from other messages, needing no special network protocols (the architecture avoids them even needing special priority - see later). Thus, although protocols wouldn't change, applications would need to hook in to on-line charging and accounting middleware where appropriate. In the case of QoS, the application only needs a QoS middleware hook adding, with the charging being called from the QoS middleware. Our QoS demonstrator [TasselBri97] has shown that both these "non-functional" aspects can be introduced in this way with just three trivial application code changes in typical cases.

The dependency of the charging system on the transmission system would not just be one way. The transmission system will probably become reliant on the charging system for its dynamic management through congestion-based price variation or initiation of capacity provision.




Fig 4 - Zero bits for charging

This offers a "zero bits for charging" solution. This is in contrast to the various discussions over how many bits are needed to indicate packets are chargeable, ranging from numbers like three [3b21] to one [Crowcroft96Kn, Clark95]. The important point here is that these transmission system bits are still needed, but for QoS (or multicast or whatever), not for charging. The contract then specifies what type of QoS (or type of other non-QoS related services like multicast) should be charged for and the mappings to locate the tariffs. This allows different ISPs to decide what to charge, what tariff structures to use and even *whether* to charge for each type of QoS at all, independently of what other ISPs decide. Thus the decoupling of the operation of measurement and payment systems from that of transmission not only offers lightweight operation, but, just as importantly, facilitates business model flexibility.

Flexibility is further assured because the architecture is based on business model generalisations that have been chosen because they support all the more specific business models [Freestone97]. For instance:

- Pre-payment is a generalisation of post-payment because the network can be operated with all the network providers internally pre-paying their own charging systems even if they offer credit

to their customers (e.g. because they wish to offer post-payment to retain customers as a business decision - along with the extra collection costs this entails, including late payers, bad debts etc.). Thus, if a BN insists on pre-payment by ENs, the ENs can still offer post-payment to their customers. Similarly, if the EN operates systems based on pre-payment for all its customers, it can offer some post-payment as a separate operation, or even through a separate finance business.

- Similarly, "pay-as-you-go" is a generalisation of batch payment, and more importantly, "account-as-you-go" is a generalisation of batch accounting. As long as the underlying technology operates on the basis of the former, the latter can be presented to the customer, but not *vice versa*.
- Spot pricing is a generalisation of longer-term pricing. It is impossible to manage a large-scale system on short-term price movements if some essential parts of the system are less dynamic. But if all parts of a system are dynamically priced, customers that require price stability can be given it (at a premium). If larger systems still are built on top of these stably priced parts (and possibly other dynamic parts), they will still function dynamically. This is because the price stability of such parts is a requirement of that part, rather than simply existing because bad design has made it impossible to provide otherwise. Clearly, price stability is often a business requirement for budgeting purposes [Barns89] and QoS stability will usually be a requirement of the application. These stability requirements have to be expressed through an expenditure controller interface discussed in [Danielsen95] and implemented in prototype in [TasselBri97]. Later discussion covers the notion that price stability itself is an element of demand which has a price.

The payment aspect of the architecture is designed around (though not reliant on) the ability to pay small amounts cheaply on the wire ("micropayments") [Hill96]. Thus it would be sensible, though not essential, for there to be a non-usage-based class of tariff for making these payments. However, the architecture is designed to allow all forms of payment technology (manual cashier-based, credit card, direct debit, barter).

It is also assumed that users will contract with network providers to run billing software on their own hosts. Again, the system is designed around (but not reliant on) the ability to manage the deployment of this software using such platform-independent solutions such as the Java run-time system [Kramer96]. However, the architecture allows any means of deployment management (floppy disk installation, software download, Active-X [ActiveX], smart card, bundled with communications stack, etc.).

Because of the above, the architecture can be applied to consumer devices that have no user interface relevant to payment (e.g. hardware Internet phones, games consoles). The payments for the device can either be enacted with a completely separate interface on another device, or payment tokens for the device to use can be transferred into it through its network interface.

More succinctly, any party may be contracted to pay for any Internet service on behalf on any other party, with any granularity at any tariffing regime by any means to any schedule.

A final generalisation is that openness can be closed but proprietoriality is much more difficult to open. Thus, the architecture identifies all the modules of the system that might be supplied by different parties. Later design work will concentrate on defining the interfaces to these and, where they are intrinsically separated across a network, the application layer protocols to allow them to inter-work without making any assumptions about a distributed object model. This is not to say that these modules represent the most granular level of objects in the system, merely the high level separations of function and state. As already stated, our earlier work in this area concentrated on a technique for simplifying the application programming interface to quality of service and charging objects using reflection [TasselBri97, Tassel97]. The assumption is that the very act of opening up the market in transmission charging software will improve the quality and value of the code created for the market.

It is assumed that any competitive market in charging software may be purely a back-end market with no apparent end-customer choice. This seems likely unless one believes end-customers will be willing to pay for software the only function of which is to bill them. An alternative market model might be that suppliers would receive royalties from the ISP dependent on how many end-customers chose their software. In this case, an ISP would probably choose to certify end-system software as compliant with the published interfaces to its charging system, rather than just publishing the

23
interfaces and assuming co-ordination. One might well expect that unaccredited software would be specified in the contract.

1.3 Price variation

1.3.1 Price variation in time

The general impression has been given so far that the architecture manages fluctuations in demand by varying price. This is not entirely accurate, so two exceptions will be described before continuing.

1. It is necessary to establish whether a demand variation is real and true. If demand is being estimated in advance, for instance by monitoring out of band "tenders" for capacity, double accounting must be avoided. Where there are multiple similar demands these may be in response to the same tender, and only one may be realised in practice. In a more sinister vein, false demand may be created by a hostile competitor in an attempt to force an upgrade in capacity, only to have the demand removed causing an excess of supply and a crash in prices and revenue compounded by the cost of investing in the capacity. This can be followed by an attack by arbitreurs, buying the spare capacity at low prices to capture any recovery in the market [{}]. Protection against such conspiracies can only be possible through commercial intelligence gathering. Therefore, the architecture merely leaves a general hook (on which to hang such business rules) in the provisioning and pricing objects.
2. As will be outlined in the justification for usage-based charging, total utility, and hence potential revenue maximisation, depends as much on the capacity of the bottleneck link as it does on price and demand. Thus, it is as important to vary the capacity of the links and routers in a network as it is to vary price. In the short term, if there is too much demand for some bottleneck resource, price can be increased to put off some of the demand and bring the operating point to revenue maximisation (see Fig 9). However, on a longer time-scale the capacity of the bottleneck has to be optimised to maximise utility and hence revenue. The art of provisioning is outside the immediate scope of this paper, so this will not be discussed further, suffice to say that it is perfectly possible to vary available capacity on very short time-scales if required, by artificially throttling actual capacity so that the throttle can be relaxed when required [{}].

Thus, price variation must not be a knee-jerk reaction. The signals that cause it must pass through the network provider's "brain" as well.

Next, as promised earlier, we must consider the fact that price stability in itself is a requirement demanded by many customers. Odlyzko [Odlyzko97] provides an excellent review of the literature lending support to this strong trend in consumer behaviour and the tendency for businesses not to offer variable pricing when they could (possibly due to expected consumer resistance). The key to solving this problem is that where there is demand, there can be a price. Just as commodity futures markets arose in response to the demand by farmers for stable incomes, so it can be for network services. In other words, the price for a certain level of service can consist of the sum of its spot price and a premium for the right to have the price itself remain stable. The price of the premium can be set to ensure there are sufficient customers buying at the spot price to be able to manage the network on very short time-scales. This band of customers will allow the provider to hold out long enough until the price offered to more stable bands can be changed. Thus, for instance, when an Internet session is announced where one party offers to pay for another, it is likely the first party will only do this at a price they know will still apply by the end of the session. However, customers who claim they need stable pricing will have a price below which they are prepared to accept volatility.

As always, there will be an element of skill and risk involved in pricing stability. One possibility available to network providers is to offer stable prices, but only with a 99% (or whatever) guarantee they won't be changed. This would allow prices to be used to manage demand during exceptional circumstances such as major outages of capacity after floods, hurricanes, earthquakes or bombings. This would be akin to having second lines of defence, but without needing to implement a secondary technological solution (e.g. bulk access control like telephone preference service [{}]). Implementing such heavyweight systems for very rare occurrences is usually extremely non-cost-effective.

Charging a premium for stability is illustrated in Figs 5a & b. The first figure shows the situation at

"time zero" where the spot price is set as shown and the shaded area represents the future possible variation of this price that the provider is sufficiently certain might happen to set more stable price on this basis. Although no scale is given (figures will depend on marketing and analysis of human psychology), we envisage the time axis being on a log scale such that there is considerable granularity of stability pricing over time-scales of minutes but also long-term price stability of the order of months is available. The stable prices are shown as price bands above the spot price, but an alternative would be for the provider to simply publish the formula for mapping the spot price to a price fixed for any specific period (its risk curve). Customers could then pick their period of stability over an infinitely variable range.




Fig 5a - Price stability flexibility: before

Fig 5b shows the same scenario in hindsight - the wiggly curve being the actual movement of the spot price with time. In particular, the figure focusses on the second most stable price. At the time this price expires, a second region of risk is drawn emanating from the spot price at that time. This enables the new value of the stable price to be calculated as shown. Typically in practice, a new stable price could be contracted to at any time, not just when the previous one had expired. However, the relationship between the spot price and a more stable price would not be simple. The more the spot price drops below its longer term average, the more the risk region is skewed upwards and this has to be taken into account in determining the formula for the stable price. How these tariffs would be announced is covered next, but only after describing one further aspect of price that will need to be taken into account.






Fig 5b - Price stability flexibility: after

Moving on to a discussion of price announcement timing, [Fig 6](#) shows that there are three events in the life-time of a price: announcement, start and end. The time a price is no longer in effect can be announced in advance or announced implicitly by the start time of a new price. Similarly, the start time can either be announced explicitly, or implied to be at the time the announcement is received.



Fig 6 - Price and time: definitions

Fig. 7 simply shows that it is possible to manage demand (to a limited degree) by judicious use of the flexibility the price announcement time offers. If demand is predicted to rise and fall in a bursty but known manner (e.g. a scheduled mass audience event), the amplitude of the fluctuation can be

26
reduced by announcing a lower price will come into effect at the end of the predicted peak.




Fig 7 - Price announcement time flexibility

If times are announced explicitly, it is necessary to define whose clock is being used. It is also necessary to ensure that neither party can stretch and expand time as measured by their clock to have longer on low prices and shorter on high. The pay-and-display model provides a possible channel for checking that clocks are synchronised sufficiently - as accounting information is regularly fed back to the provider, the current time could occasionally be reported both ways to ensure any cheating is confined to the maximum deviation of the round-trip time between customer and ISP.

If the times of price changes are implicitly defined, they cannot, in all fairness, come into effect at the time they are sent by the provider. They can only be defined as effective at the instant of receipt (and could therefore be different for each customer). Our work on securely encapsulated key generators [Fairman98] provides a relatively lightweight way to prove the time of receipt of a multicast message without per message acknowledgements. However it relies on use of a smart card or secure co-processor, and the clock of such devices has not been the subject of any concerted tampering or tamper-proofing efforts to date. While the time between price changes is large compared to round-trip times, such solutions will probably be unnecessary. For instance, if pay-and-display accounting is used, as long as the price in current use is declared in the accounting information, it should be possible for the network provider to detect any fraud outside the deviation of normal round trip times as described above.

It should be possible to openly publish all prices and even the algorithms used to derive them from each other. This stems from the fact that pricing is the way a provider announces how it wishes to co-operate with its customers. However, it is more than likely that there will be a business requirement to hide pricing for certain sectors from others. If, as proposed, pricing is to be announced using receiver initiated multicast, this would have to be achieved by encrypting the pricing channels. The problem then becomes one of managing keys. Our work [Fairman98] solves the problem of doing this with minimal messaging (zero effect on other users when anyone joins or leaves). Other approaches are described under Related Work.

1.3.1 Price variation with localised load

So far price variations have been discussed as if everyone (at least every customer of a particular ISP) will be operating on the same price for the same service class. The mechanisms so far described vary the price for all customers so that the differential between levels of quality can be maintained given fluctuations in demand for each level as a whole, regardless of source or destination address.

We must also avoid end-systems getting involved in least cost routing (see What Business earlier)

24

we must also avoid end-systems getting involved in least cost routing (see What Business earlier) unless we wish to have the network and the end-systems entering into a perverse demarcation dispute over who should decide the "best" route. However, we do want end-systems to be able to choose the least cost remote host offering a service (regardless of route) or the least cost time when to use the network (if at all) or the least cost level of quality to achieve their ends. As outlined earlier, as long as we trust the network to provide us with the best route between any two points, network routing can be encapsulated and stay outside the concern of hosts.

Thus, where there is congestion, the first response should be to re-route around it. If it is not general congestion, but just congestion for a particular level or type of QoS this is the domain of work on QoS routing, a can of worms just re-opened by the IETF [{}]. This is outside the scope of this paper.

However, once any available re-routing has been achieved (or the attempt has been given up as this is often difficult in practice) we are still left with the need to deal with congestion. As justified later, ideally we would like to do this with price-based rather than "last in first barred" admission control.

If we rely on the multicast channels described earlier to signal price changes, it would be extremely difficult to cause a price rise signal to be multicast out in such a way that only those wanting to pass through the congested point would join the multicast. Congestion at the other side of the Internet is just as important as congestion this side, if both bottlenecks are on the required path through the Internet. There would either have to be multicasts for every element in the Internet (e.g. every queue on every interface of every router) or some form of aggregation would be necessary (as described by Shenker in the edge pricing proposal [Shenker96]). However, aggregation would mean that many users not passing through the congested resource would be given price rise signals. This implies the price rise signals would have to be damped by the weighting that the non-congested resources had over the congested one. This would lead to insufficient demand back-off from congestion leading to an inability to provide the quality requested even though a higher price were being charged - not a recipe for customer satisfaction.

Fortunately there is another way. Variation of price with congestion can be triggered by any of the congestion control mechanisms already in place or being developed. In all cases, by design, congestion control signals are detectable by the end-systems causing the congestion and only those systems. Thus all that is necessary is to define up front (with the tariff) how the price will vary with congestion. An algorithm mapping congestion level to price can be installed in the host through the same channels used for the tariffs, contracts etc. Then, when a congestion signal is received (even implicitly as in the case of packet drop), the host can calculate the matching price rise for continuing to utilise that path through the Internet. Thus *price* changes (as opposed to *congestion* changes) can be communicated with zero bandwidth during operation.

Through the application of user policy, the end-system can then decide how to react to the price change, rather than being coerced into a reaction to congestion specified in a protocol. If the communication path is important to the user, they can simply accept the price rise and continue at the current utilisation rate. If everyone adopts this approach, the congestion will continue to worsen and the price will continue to rise. Eventually the point will be reached where enough users have been put off by the price rise for the congestion to cease. The algorithms will have to be designed to ensure the price rises steeply if congestion persists while also ensuring the initial price rise is low enough not to cause oscillatory behaviour. As there are no extra round trips (because the price sensitivity is internal to the end-system) the response time will be unchanged from the unmodified congestion control schemes (except for some minor extra host processing). If a delay (e.g. user intervention) was involved in the decision to continue at a new price, the transmission rate would have to react as for the unmodified congestion control scheme if transmission were to continue in parallel while waiting for a decision. The same scenario would apply if the signalling was of impending rather than actual congestion.

Some current congestion control mechanisms (e.g. TCP) hide the distinction between network congestion and remote end-system congestion. For efficient signalling the required congestion window can simply be the minimum of the two. However, if load signalling is to be used to potentially control the price paid to the network, the two must be distinct. Many congestion control protocols are still in the process of definition so this distinction can be made. We have no such freedom with TCP, but an approach that uses some of TCP's redundancy is described later. In the final analysis, if there is obviously no change in the capacity available when an end system continues to ignore congestion signalling, it can infer that the congestion is not due to other more yielding users

and conclude that it has to back. The congestion might, for instance, be on one of the access links with no other users responsible.

Recalling the discussion of demand for stability earlier, those users who are paying a premium for price stability would have this reflected in the algorithm installed in their host connecting congestion signals to price changes. In other words, they wouldn't see a price change for as long as they had paid not to see one. The ISP would have to weight the algorithm of those users *not* paying for stability accordingly, so that the price rose (or fell) more steeply than if all users prices were reacting. For this to work, there would have to be enough acceptance of slightly bursty prices in the market to ensure sufficient proportion of the possibly small number of users through any one bottleneck would exist to make the network manageable. If the price differential between stable and volatile pricing had to be too great, or the level of fines for ignoring congestion had to be too high, it might distort the general structure of tariffs for other services. However, some of the oddities of human behaviour [Odlvzko97] might be smoothed out by the general introduction of agents working more rationally on behalf of users.

It sounds dangerous to vary the price with network loading in such an automatic way - only a few paragraphs earlier it was said that price rises should not be a knee-jerk reaction. However, this is exactly what is being proposed - a locally controlled reflex. However it should be noted that this is only being proposed where (imminent) congestion signalling is itself a method of last resort. This is because we are talking about congestion of a single differentiated service, not of all services through a bottleneck. Prior to that, the following actions, as already discussed, would have had to be taken to the limit of their effectiveness:

1. the network re-routing around congestion
2. the network borrowing capacity from "lower" levels of service (lower in the context of the relevant dimension(s) of QoS) including the best effort service
3. the network introducing extra capacity (possibly automatically)
4. the end-system establishing that the congestion is on the shared network and not just on the access links or end systems
5. the end-system setting QoS requirements to a "higher" level (if cheaper than the fine for ignoring congestion at the current level)
6. the end-system deciding it is essential to ignore the congestion, given the fine for doing so might be quite high
7. both (all) end-systems agreeing to ignore the congestion

In summary, by introducing a level of indirection in congestion control mechanisms we have achieved price-based congestion avoidance or control without any need to change existing protocols. Essentially we have the ability to calculate exactly, the complete range of prices for permission to contravene traditional congestion control rules to any degree so that anyone can choose to break the rules at any time and know the penalty. At the same time, with pay-and-display and sampled policing, providers don't have to take the hit of policing whether customers are paying the correct price, except on an occasional basis. A discussion of the finer detail of how these price rises would be policed, and other knock-on effects is left for later.

However, changing the congestion control response of the Internet has to be treated with great care. The predominant algorithm of additive increase and multiplicative decrease (TCP) is very stable and we are proposing to replace it with a highly heterogeneous response, although it should be clear that this is a response of last resort. This will require simulation and market trials to determine the best algorithms

2. Justification

2.1 Justification for a usage-based charging sector

So far this paper has followed a rather unorthodox order; diving straight in to a high level description of the proposed architecture without discussing the requirements that led to it. This was deliberate. All business model proposals that become accepted exploit what is practical. The mentality that says that telecommunications services must be charged by measurement and billing enacted by network providers alone is so entrenched that it would have been difficult to write an understandable paper without exploring the expansion of "what is practical" first. Having done this, we can now move on to

ify the approach more carefully. We start by justifying usage-based charging itself then continue a justification of the architecture proposed in order to achieve this efficiently.

The Internet Protocol and the Transmission Control Protocol (TCP/IP) [Thomas96] have tempted millions of computers across the globe to get connected to the Internet. This is because it has proved possible to build thousands of distributed applications on the foundations of these protocols. Such a wide range of actual and potential applications combined with such a large number of actual and potential connected end-systems reached critical mass in the early to mid-1990s. It became difficult for anybody to believe any other technology would form the global information infrastructure.

All this was achieved with essentially just two business models^(iv): flat-fee Internet access subscription charges proportional to maximum access bandwidth and some operators levying additional time-related connection charges. In addition, many ISPs choose (in many markets there is no choice) to offer service over access links that have time-related charges (or sometimes traffic related charges in the case of inter-ISP links) levied by the telecommunications operator. When there is no congestion in the core network, end-systems are only limited by the lesser of their own capabilities and the size of the access pipe that has been purchased. The incentive to buy a bigger pipe has been effective only because there usually isn't congestion downstream of the access on a sufficient proportion of routes due to considerable fan-out, usually leaving at least one destination of interest clear of congestion.

One of the subtle elegances of TCP that has fed Internet growth is that it defines a fair share of a congested network and a lightweight mechanism to achieve it [RFC793, RFC2001, Jacobsen88]. It maintains fairness by setting up an implicit signalling channel from each end system to the worst bottleneck on the network path. All the bottleneck has to do is throw away packets during congestion and the end-systems then co-operate to adjust their rate to tend towards what TCP defines as their fair share, taking into account the dynamics of all the other relevant data flows. The importance of this was that demand management was achieved without usage-based charging. This allowed the Internet to develop without having to wait for a dynamic usage-based charging infrastructure. The business model based on this definition of fairness has proved particularly appropriate both within non-commercial organisations and, more surprisingly, within and between large or small commercial organisations, whether or not they operate internal markets to manage their demand for internal services. Thus, preservation of this sector of the Internet that can survive on simple flat-fee access charges is seen as paramount because its running costs are minimal.

It must be acknowledged that the early TCP didn't have congestion control. TCP evolved when the bottleneck moved from the receiver's capacity to consume data to the network's capacity to forward it, which was a consequence of the explosive increase in numbers of end-systems. Similarly the random early detection (RED) [Floyd93, Braden97] proposal evolved due to a worrying increase in the amount of end-system software that either didn't adapt at all to congestion or adapted more greedily (by accident or design) than the TCP norm. Because the network could no longer trust end-systems to behave fairly, RED-enabled routers are being introduced to ensure greedy applications can gain no advantage. This is achieved by discarding packets randomly from within the router's queue during congestion rather than dropping the tail of the queue. This has a detrimental effect proportional to greediness rather than allowing greedy flows to hog the queue to the detriment of responsible flows.

Thus a definition of fairness has evolved that relies on a combination of co-operation and policing between the end-systems and the network. This definition can most strikingly be paraphrased as "to each *proportional* to their need, from each according to their ability". Ironically, in non-capitalist parts of the globe, the Worldwide Web, which is the most popular application built over TCP/IP, is epitomised as the tool and product of the capitalist system [Xiaoming96], with its unquestioning assumption of a *laissez-faire* economic model^(xiii). This irony merely demonstrates how the distinction between capitalism and communism depends on the definition of "need". On the Internet the "need of each" has been clearly defined and scoped and it can now be policed.

Herein lies the problem. A whole range of applications don't work reliably on the current Internet because their needs are greater than what they typically get allocated under this definition of fairness given the current demand, supply and price of Internet capacity. These are the real-time or inelastic applications like Internet telephony or video conferencing, which need minimum resources in order to operate at anything approaching levels people are prepared to endure even when paying nothing more than for their own time.

Shenker [Shenker95] provides a now classic distinction between non-elastic and elastic applications (as well as covering very similar ground to that in this section). However, it is necessary to digress from the thread of our paper at this point to amend the definition of this distinction before continuing.

The theory goes that although users of elastic applications derive decreasing utility from a congested network as the number of users sharing it increases. However, the total of all the marginal utility reductions of all users is less than the utility that each new user gains on entry - the utility curve is concave. The graph in Fig 8a shows Shenker's *Gedanken* experiment for a single link shared by a number of users of one elastic application. It shows the utility that one user experiences against the responsiveness ($1/\text{delay}$) of the end to end network path (utility is always assumed never to decrease too, which is reasonable). The horizontal axis can be any aspect of the service consumed - Shenker used bandwidth in his example. The horizontal axis can also be considered as the inverse of the responsiveness of the bottleneck reducing as the number fighting over it increases. The characteristic of an elastic application is that some utility is derived from even the smallest responsiveness (otherwise it wouldn't remain concave). Thus, as n increases, $1/n$ approaches zero but there is still finite utility for each user.

In contrast, Fig 8b shows a similar graph for one inelastic application where there is a point where late arrival of information is no better than non-arrival. Thus as we move to the left as more users share the congested resource, everyone starts to lose all their utility at once as the utility function of each user collapses long before $1/n$ approaches zero (i.e. still with a finite number of users). Moving to the right (less users sharing leading to more responsiveness) of the steep part of the curve, we enter the region where more responsiveness is available than the application can usefully exploit. This is why the utility curve flattens out. An example might be an Internet phone where it becomes less and less important to the customer whether the delay in hearing the replies of the other party is a few milliseconds or a few nanoseconds as long as it is not a few seconds. Elastic applications enter a similar regime where the network is effectively over-dimensioned.

Note that the horizontal axis may represent another resource that is scarce for the particular application in question, such as bandwidth or reliability, but for many applications responsiveness is the resource in critical shortage before others.




Fig 8a - "Elastic" application






Fig 8b - Inelastic application

From a network provider's point of view, the goal is to dimension the network such that for most of the time the number of customers through a bottleneck balances the quantity of bottleneck resource at a level such that the total customer utility is maximised. Customer utility represents the maximum price a customer would be prepared to pay, so total utility represents total potential revenue. Fig 9 shows this graphically for an inelastic application where the rightmost vertical represents the bottleneck resource for an individual on an end to end path. As more customers share the bottleneck, the responsiveness available to each is progressively shared into smaller fractions. But, the utility each user gets from this responsiveness must be multiplied by the number of customers in order to derive the total utility. For this inelastic application there is clearly a peak in total utility (near the turning point of the utility curve), which implies the price per customer should be set at the utility that individual customers derive under this peak, u_0 (or just above it, u_q , if the ISP wishes to exploit quantisation of utilisation). In practice different people have different utility curves, but average customers are sufficient to explain the principles (see [Kelly?]) for a statistical analysis).




Fig 9 - Inelastic applications sharing bottleneck

When the equivalent formulae for total utility are derived for the family of elastic curves like that in Fig 8a, with increasing customer numbers one get a rather odd result. As the congested resource share of each customer tends to zero the total utility summed across all customers tends to a non-zero value. This leads by a *reductio ad absurdum* argument to the conclusion that, in fact, there is no such thing as an elastic application^(xiv). This is supported by intuition; it seems fair to say that every communications application needs some minimum level of resources to be worth using at all. For instance, e-mail is still usable with responsiveness of days^(xv), but beyond that it hits its limit as more and more messages arrive after they are no longer relevant. So-called elastic applications are in fact *inelastic* but with the operating point very close to the zero resource axis (as shown by the alternative graph (dotted) overlaid on Fig 10). The distinction between elastic and inelastic isn't even related to whether the responsiveness is of the same order as the network half-round trip time, because the RTT is only absolute for an empty network, otherwise depending on the economic environment.




Fig 10 - "Elastic" application - no such thing

So why is this academic point scoring relevant here? The important point is that there is no qualitative difference between applications, only a quantitative one. What in fact is going on here is illustrated in Fig 11. Here we show the utility a single user derives from an end-end path against level of responsiveness when considering all applications each user may have. Because the application space is not continuous in its requirements, the total utility curve has definite steps as the operating point for each "killer application" is passed. The curve for total utility is only approximately the sum of the curves for the killer applications because other less prevalent applications would contribute to its total, probably resulting in a slight smoothing as well as a general lifting of the steps. Without differentiated service classes, as is the situation today, at a certain average flat-fee for access a certain average number of users will be sharing a bottleneck at any one time. Thus although all applications are inelastic, applications to the left of this operating point seem to be elastic. This is because access charges are a barrier to the market, preventing us ever seeing the ridiculous "near-infinite" numbers of users all sharing one bottleneck in the *reductio ad absurdum* argument above.




Fig 11 - Multi-application network

reason for the digression should now become clear. It appears inevitable that usage-based charging should be introduced to allow access to the more demanding requirements that are currently above the economic balance point that is the maximum that can be achieved with flat rate charging. However, before jumping to conclusions, let us consider how we might still avoid usage-based charging.

In order to avoid usage-based charging one might try to evolve the Internet protocols so that proportional fairness could be stretched to encompass the tighter deadlines or heavier resource requirements of applications currently too inelastic to coexist with more elastic applications. One might try to evolve the network's policing capabilities to ensure that more elastic software wasn't claiming inelastic needs. However, TCP/IP appears to offer a remarkable compromise for where we draw the "line in the sand" of fairness. Trying to embrace such tight deadline-based needs would probably break this consensus. If the so called real-time applications (those to the right of the current operating point) represented only a small part of the potential market, this might be a feasible proposition; people will allow a person who is genuinely desperate to jump a queue only so long as the number of desperate people is small. This is almost certainly not the case with real-time applications.

An approach considered by some ISPs (and desired by many managers of corporate networks) is to use policy-driven allocation of resources rather than charging. Packet shaping boxes offer tempting capabilities such as giving priority to packets containing the TCP protocol if addressed to or from the corporate MIS system, or to prioritise IBM's SNA protocol over IP because it is most used between mainframes. Such policies are only practical at the edges of networks. They quickly become impractical at large scale, where the single dimension of price appears to be a good common denominator to determine demand and to cover the cost of that demand. The limitation of a pricing mechanism is that it is bad at expressing absolute policies like "no access to porn sites", which move into the realm of security. Further work is proposed to see whether all such policies can be modelled by a market mechanism. This might be based on the allocation of tokens which behave like money but which have no direct conversion to a currency of the world's monetary system because they can have on-off value as well as variable value.

Even if there is unlikely to be a protocol or policy-based solution that will allow the co-existence of such a broad spectrum of applications on the Internet (as currently dimensioned), there might still be no need for usage-based charging. It might still be feasible for the operating point to simply move further to the right in response to the market. After all, the operating point (where responsiveness is concerned) managed to move from 'e-mail only' to 'e-mail plus Web' because of the size of the Web market (e-mail is now running on an over-dimensioned network by cross-subsidy from the Web).

This is only feasible as long as the demand for the next killer application is sufficiently bigger than the demand for the previous one to make the total revenue curve peak at this next step with such a size that it absorbs the cost of over-dimensioning for the Web. It would appear that telephony is already a "killer application" of sufficient size to motivate this. However, the rather important factor against this happening is that there is already an alternative network that provides this service which uses usage-based charging rather than over-dimensioning and is therefore going to be competitive as long as the cost of the usage-based charging system for telephony is less than the cost of over-dimensioning the Internet. In other words, the question is, if every customer of an ISP paid a flat fee subscription for their Internet connection equal to about as much as their combined Internet and telephony bills, would this provide sufficient cash to over-dimension that ISP's network such that it could run a decent telephony service over it without usage-based charging? Even if this were the case, it still leaves the network under-dimensioned for yet more demanding applications^(xvi). This raises the question of whether there will have to be a usage-based charging system anyway for these applications, so the cost of it could be spread over less stringent requirements instead of investing in over-dimensioning. Note that such decisions can be taken on a per-ISP basis - inelastic applications should still work across a mixture of over-dimensioned and usage charged networks.

There is a third alternative that might yet avoid usage-based charging (if the scenario in the previous paragraph doesn't materialise). This is to charge a series of incremental flat fees for unlimited access to each congested resource up to incremental maxima, which match the killer application bands. The Paris metro pricing (PMP) scheme is an example [Odlyzko97]. Given the known large market for telephony, this might well reduce to the large majority subscribing to all the bands except for the top one dedicated to network gamers. Thus this would probably become just a slight variation on the

34

re dedicated to network games. Thus this would probably become just a slight variation on the previous scenario (except with a lower barrier to entry at the bottom end of the market). The difference is that it wouldn't be possible to differentiate between telephony customers who didn't use the Web and those that did, because access to a higher level would automatically open up all the lower levels. Another problem is that communications takes place between multiple parties, not just one, so everyone would have to be on the same logical network to communicate. In other words, the solution doesn't allow independent commercial decisions. Also, the important similarity to the previous scheme would remain that it would require over-dimensioning for each band. A variation would be to allow customers to just pay for the time they need access to a higher band but impose a lower threshold on this time to prevent signalling overload of the charging system due to flapping between bands. This falls within the definition of usage-based charging (and could be achieved with the proposed architecture) but it has the additional benefit that it keeps the number of chargeable events down. However, it appears to encourage very bursty use of the network and consequently will need further analysis and simulation (see Further work).

In fact, any flat fee system is in imminent danger anyway, from a more widespread adoption of goal-driven agents (e.g. some future incarnation of a robot like WebWhacker [BlueSquirrel]). An obvious rule to add to the goals of such an agent would be to utilise to the full any flat-fee payment. This even endangers the traditional ISP business where profits are made to a large extent on the difference between revenues from customer access links and the cost of backbone access links. With humans as the limiting factor on bandwidth utilisation, the sum of all the potential customer access bandwidths less the typical proportion of inter-customer traffic is orders of magnitude higher than the actual back-bone bandwidth needed because customers' use is sporadic. If everyone runs agents which permanently fill their access pipe, profit is more reliant on the proportion of inter-customer traffic which can only be increased by attracting a greater customer base.

To be absolutely clear, we must repeat that we wish to preserve the "elastic" (less inelastic) applications sector of the Internet that so successfully avoids usage-based charging. In adding a usage-based charging sector, the aim is to complement the flat-fee charging sector not to repress the current fragile economy of the Internet and the communities that have appeared on it. Adding usage-based charging will cause flat basic access fees to drop thus reducing the barrier to entry into the Internet market, growing the total community. It should then be possible for the Internet to continue to appear to be an egalitarian community free at the point of use, while simultaneously supporting commercial applications and bringing in the revenue to pay for the necessary infrastructure evolution and growth. "This would leave undisturbed the cultural aspects of the current best-effort Internet..." to quote from a passage in Shenker [Shenker95] which echoes an assertion we probably all try to believe - that we will not be guilty of deliberate "cybercide". A pessimistic view could predict that, once a charging infrastructure is in place, it will be cheaper (than over-provisioning) to use it for all classes of service. However, charging costs are, in the main, running costs. Therefore, it still seems likely that the base best-effort service will remain flat-fee charged. Indeed, one of the few pre-requisites for the proposed architecture is a reasonably large proportion of Internet traffic remaining best-effort flat-fee access. It also makes good business sense not to charge more the longer your customers look through the shop window, even if they pay taxes for the pavement.

As a subtle example of what might happen, an infrastructure such as that proposed could be used to charge differentially for multicast over unicast, despite them both being best-effort and both having nearly identical underlying costs. In other words, the only reason there is pressure to charge a premium for sending to a multicast is that a profit can be made until the market becomes competitive. Thus, if this infrastructure were used to charge a premium, the "cybercide" could be construed to have started.

It is outside the scope of this paper to attempt to answer the market prediction questions that we have identified will determine whether usage-based charging is needed. The main conclusion so far is therefore simply that the need for the proposed architecture depends on whether it can make the cost of usage-based charging less than the cost of over-dimensioning. This motivates more detailed work to enable estimation of the cost of various usage-based charging proposals, in particular generic ones such as that under discussion here. This is made more difficult by the fact that the proposals each live in different markets so are not directly comparable:

- over-dimensioning (network capacity market)
- flat-fee charging with banded over-dimensioning (hybrid of charging software and network

- capacity markets)
- usage-based charging (ISP boundaries (charging software and measurement software markets)
- accurate cost plus margin charging, for an example see Herzog *et al* [Herzog95] (new accounting mechanisms markets plus charging software markets)

So although we haven't justified usage-based charging to the point of incontrovertible proof, we must now move on to justify the approach we have proposed to achieve cheap, usage-based charging in more specific terms than already given.

2.2 Justification for approach

The proposed architecture has the following novel aspects, which each need justification:

- pay-and-display model, which is justified under the following sub-headings:
 - scalability
 - non-repudiation
 - measurement accuracy and relevance
 - accounting dynamics
- pay-as-you-go
- price-based access control
- independence between charging and transmission systems
- independence of network layer business model from higher-level business models

All these features are argued to be valid improvements in their own right, independent of each other, but together they form a coherent architecture too.

2.2.1 Pay-and-display justification

Scalability

Firstly, in order to operate an ISP business, the priority must be to collect revenue at minimal cost. The revenue collected from each customer doesn't have to 100% accurately reflect their use of the services as long as they are happy and other customers don't have grounds to complain of unfairness. An obvious way of achieving this is to accept that a chargeable event may occasionally be missed and build an allowance into pricing to cover this (cf. the economics of the corner shop, with respect to spillage, pilfering and perishing). The concept of the customer billing themselves with random checks is common where very low price items are being traded, e.g. pay and display or metered car-parking, or self-check-out in certain supermarkets (since 1996) [U-Scan]. The principal behind all pay-and-display systems is to utilise a tiny proportion of the customer's resources to do charging, rather than having to explicitly spend centrally which may have economies of scale but can't be concealed from the market. As there will never be dumb Internet terminal equipment that cannot run software, the opportunity presents itself to bill in "middleware" on the customers' systems. This appears to swap a scalability problem for a configuration management problem, but the architecture relies on software deployment techniques that are becoming more commonplace (though still not fully tried and tested) [Faupe197] to contain this.

In general, traffic measurement systems are designed for traffic analysis, not settlement, although some could be used for both but wouldn't necessarily be scalable [Carpenter96]. Losses of analysis traffic can generally be catered for statistically. If allowance for loss of settlement traffic is to be designed in, one opens up the possibility of fraudulent deliberate loss and claims of unfairness between customers. Already the view that the only practical way to measure heavy packet flows is with an unreliable protocol has become hegemony. The IETF are basing standardised flow measurement [RFC2063] on simple network management protocol (SNMP) [RFC1905] because it in turn is based on unreliable user datagram protocol (UDP) [RFC1906, RFC768]. No one is proposing measurement protocols should add any higher level reliability of their own either. This is because, if reliable measurements were to be implemented, there would be a consequent build up of state (for retransmission buffers and control) in the routing systems with an inherent positive feedback problem during network congestion (the alternative of over-dimensioning the measurement system appears to exhibit prohibitive cost). A scheme that is more likely to lose accounting information just when it needs it most (to collect revenue when prices and volumes are likely to be at their highest) is not particularly interesting. Note that bulk measurement is not prone to the same problems. So avoiding



company that is trustworthy as one that isn't. Specifically, if one ISP's prices are cheaper than

It should be noted that pay-and-display doesn't give cast-iron non-repudiation, but the improvement it does give comes for free along with the other scalability benefits discussed above (not strictly free - there are the deployment and extra messaging costs). Where a system doesn't have non-repudiation built in, there are numerous fraud possibilities which are rarely obvious to the designers. Where it is built in, even if not perfect, the fraud modes are explicit. Traditionally, with billing done in the network, the system design is kept private. It is generally agreed that security by obscurity is a dangerous strategy as internal or ex-employee fraud is the biggest threat (and probably goes mostly

34

measured). As combating fraud is an ever outward-spiralling problem, it is best that control over level of trust in each user can be explicitly set, rather than just hoping no one will notice the holes. With this architecture it is possible to sample different users with different frequencies and even to supply them with different accounting software, possibly without their knowledge.

Finally, it is expected that psychologically (and rather intangibly), customers will put far greater trust in measurements taken on their own system, particularly if the software is independently accredited. This may also help improve acceptance in the Internet community among those opposed to usage-based charging, particularly if it is clear that basic access prices will fall when it is introduced.

Measurement accuracy and relevance

Ideally, there would be a "measurement region" at the interface between network provider and customer within which it could be guaranteed that nothing being measured would change as it passed through the region. Within that ideal region would be two measuring systems, one trusted by the customer and one trusted by the network provider (a-a in Fig 12), or alternatively one system trusted by both (b).

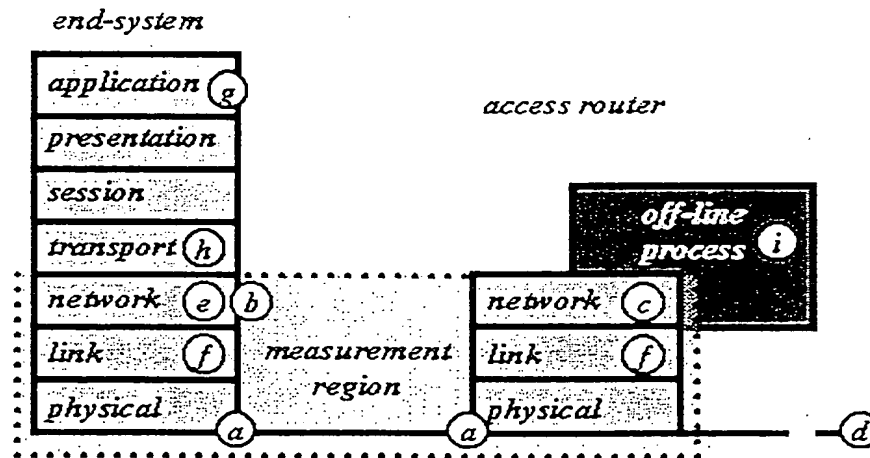


Fig 12 - Measurement regions

A reasonably lightweight (b-type) solution that creates such a region trusted by both parties (on a smart card connected to the customer's end-system) is described in [Fairman98]. This appears impractical for network level charging, as it involves users operating smart cards for encryption (with keys owned by the network provider(s)) whether or not the users themselves need the information to be encrypted. If encryption is to be avoided, the problem with creating a measurement region trusted by both parties is that it has to be open to physical inspection by both parties (or a party trusted by both). Otherwise either the customer or the provider can alter the topology of the measurement system to respectively either make some data bypass it, or to inject spoof data for measurement then extract it later. Given that general purpose processors are already being operated by the customer and the network provider to do the communications, it would seem efficient to use these for measurement as well, rather than having to provide special measurement devices in trusted interface regions. Thus it seems that we have to relax the ideal requirements and allow for two measurement systems with a possibility of some discrepancy between their results.

For the network provider, the practical measurement point is relatively easy to locate - at the first access device(s) for each customer that inspects network layer headers (c)(IP in this case). ISPs should not measure any deeper into their network (d) because their access network and systems will introduce delays and losses. In the Internet model, any realisation of contracts with the customer would appear (or be referred to through the higher layer protocol id) in the packet header at the network layer [RFC791, IPv6]. Because the Internet architecture aims to provide "IP over everything"

network layer [RFC791, IPv6]. Because the Internet architecture aims to provide "IP over everything" and because this has all but happened in practice, measurement at the IP layer also reduces the number of different classes of measurement software needed.

For an individual customer (e.g. on dial-up access), a practical point at which to measure would also be alongside the network layer but in their end-system's stack (e).

Ideally these measurement points would be lower in each stack to be closer to the interface between the two parties and less likely to be affected by contention in the stack. However, measuring at the link layer (f-f) would be inappropriate because only some chargeable parameters set at the network layer will ever be reflected in link layer frames; network level multicast, end-end latency requirements etc. may never be visible at the link layer. Also, of course, link layer headers would need to be ignored when measuring packet sizes for bandwidth calculations to avoid apparent discrepancies where different link technologies are chained together. In other words, the network layer is the layer at which an ISP is offering the business under discussion.

In the reception direction (up the stack) this implies that the lower layers must be dimensioned (buffer sizes, interrupt and thread scheduling priorities) to cope with the most stringent QoS requirements of higher layers. As frames are taken off the physical media, the machine must be able to pass data up the stack without any chance that usage-charged data gets discarded (e.g. due to buffer overflow caused by interrupt contention) before it gets to the network layer. It is at the network layer where the ISP's service is to be measured and where it is most convenient for QoS requirements to control correct differential treatment of the various flows as they are passed further up the stack (on end-systems) or forwarded (on routers). As this is a general requirement on the lower layers, not just for charging, such dimensioning should end up being the case in practice. This not only applies to contention in the stack, but also on shared physical media (for transmission in both directions), which is indeed being addressed by such work as the IEEE802.1p draft [IEEE802.1p] for defining QoS for all the IEEE802.x standards, such as the Ethernet family.

Whether the customer's requirements have really been met can only truly be measured by the application being used. However, it is inappropriate for the customer to measure what they expect to be charged for at the application layer (g) because any losses or delays may have been caused by contention above the network layer in the stack on their own machine.

The possibility of selling the right to vary the implicit contract that the TCP standard imposes on users was introduced above. This involves measuring the rate of increase (and decrease) of the send or receive window size. This is relatively easy in the transport layer of an end-system (h) but requires considerable off-line traffic analysis at a router (i), as this rate can only be derived indirectly by isolating the various flows apparent at the router and monitoring traffic growth over time. This is therefore impractical whether traditional network-based billing or pay-and-display is used, because even sampled network measurement would be too expensive at scale. An alternative might be to control something similar to the cisco proprietary weighted RED [Cisco] to make it drop less packets from flows that had paid for the privilege. As already explained, unless the router acts on per packet signalling bits, it has to identify transport layer flows heuristically, which is done at (i) and is therefore processor intensive. For instance, experience with flow RED (FRED) [Lin97] showed it took some time to detect a flow after it had first started so this is unlikely to ever be a scalable solution. UCL's MultiTCP [Crowcroft98] solved this by considering a single bottleneck link and concentrating on measuring the TCP window rate on a Web proxy at one end - an "end-system in the middle". This is obviously impractical as a general solution for flexible, high performance Internet charging.

Thus, for a customer on an individual link, the ideal measurement region would be the extent of the link between customer and ISP, but it can be extended to encompass the stacks up to the network layer at either end (c-e). In this case customer measurements would be no less accurate or relevant than ISP ones. Indeed, customer measurement could be less inaccurate in times of congestion when considering the measurement system as a whole. In the end, the interface between customer and provider for charging purposes can be defined as wherever it is agreed to put the meter. This mirrors the water industry regulations in the UK, in which the water meter defines the point where the responsibility for leakage changes hands, even if the meter is not anywhere near the boundary of the property.

For customers with numerous end-systems on their own network, which is in turn connected to their

39

pay-and-display by all end-systems *seems* less sensible. However, there are many special cases where it would be reasonable to include these complete networks in the measurement region. A typical example would be a small business or domestic network where one could assume network contention was low due to over-dimensioning. Such special cases can be extended to any over-dimensioned network, which opens up the interesting possibility that it could also apply to each logical network "band" carrying non-best effort traffic in most typical network QoS schemes (except controlled loss). If, as seems sensible and likely, these bands sit on top of each other and ultimately on top of a band of best effort service, they can all be considered to be over-dimensioned. This is because they can eat in to the band below and ultimately eat in to the best effort "buffer" band whenever necessary. In other words, even for numerous corporate end-systems, collated measurements of high-QoS traffic destined for non-internal addresses can be no different from a measurement taken at the egress link. This opens the possibility of inter-ISP accounting being based solely on end-system measurements and ISP sampling. This possibility is explored further in the discussion of the accounting and payment architecture. The discussion considers such scenarios as an edge network with multiple access links to backbone networks or other edge networks. In order for a network operator to split its accounting data between these multiple possible network providers, it would need to reference its link-state routing table, or at least a moving average of it. This would have to include multicast routing so that it could calculate whether a multicast tree split within its own network or not, to avoid double or under accounting.

The ideal solution that introduced this section specified that nothing being measured would change as it passed through the region. Just as ideal for pay-and-display would be two inaccurate measurements where each was equally as likely to be high as low by the same amounts. This can be achieved, for instance, when catering for system crashes - any measurement in the process of being recorded to stable storage can be designed to have a 50/50 chance of being either over or under measured by the same amount after a re-start.

Accounting dynamics

A particularly strong reason that justifies pay-and-display is that the owner of the end-system is likely to need to know the current state of their accounts more often than is the network provider. As systems develop that need totally dynamic accounting information, the network provider will be able to use bulk measurements for all foreseeable purposes - there will be no need for an up-to-the-minute view to be built bottom-up by adding together every customers' accounts. Rather it would be built top-down from indicators working over longer time-scales. For instance, dynamic pricing would be achieved by measuring bulk queue sizes and new tariffing policies would not require detailed knowledge of the status of every account in order to calculate an aggregate.

However, each customer will be running systems that make decisions based on the exact current level of their own account and on exact comparisons of the price of various alternatives.


Thus, if the accounting originates at the end-system, it can efficiently form the basis of new higher-level customer business processes. It only needs to notify the network provider of the account status at relatively frequent intervals (e.g. to allow policing). If the accounting were done in the network, more frequent messages would have to be transmitted to keep the end-system informed than if done the other way round.

Pay-and-display justification summary

The primary motivation for using pay-and-display is scalability for a given cost. In moving to measurement at the end-system for individual links (e.g. dial-up) the accuracy of measurement is (theoretically) not impaired and, if anything, improved under congestion conditions. It is also possible in theory, though by no means certain in practice, that the pay-and-display model could be operated on whole networks of end-systems as well as single system customers. Moving to pay-and-display has the added benefit of non-repudiation, which although not currently *perceived* as a problem could become a major issue overnight. Lack of non-repudiation, however, is, in fact, already a problem for communications providers in that it carries a large cost in terms of customer bill query support. Finally, pay-and-display puts the primary source of accounting information closer to where it will be needed more frequently - on the end-system. The network provider will need individual account information less frequently because dynamic processes will rely on bulk measurements, not bottom-up aggregation from individual accounts.

tion

every interested customer is contrasted to the network signalling congestion through a price rise






Fig 13 - Admission control at source

As was outlined earlier, multicasting tariffs can only manage generalised congestion of each service class for large groups of customers. To deal with more specific congestion, the alternative of notifying the end-systems in advance of the appropriate price response to localised congestion is justified only where more general management through pricing has failed. This falls back to the signalling model of the upper half of Fig 13. However, each service only requires congestion *signalling* - the complexity of *admission control mechanisms* can still be avoided, as admission control is still achieved at source.

Price-based access control appears to run counter to the requirement of many customers who want price stability so that they can budget effectively. However, we have already discussed how it is possible to offer price stability while simultaneously managing demand with spot pricing! The key to the trick was essentially to sit customers in bands depending on how stable they want their pricing to be and to manage short term congestion with the volatile bands until longer term congestion can be averted by changing the prices of the more stable bands.

Also price-based access control is based on the assumption that customers are price sensitive. One might think that, where a third party is paying (e.g. the end-user's company), this might not be the case. Two points counter this argument. The first is that, to manage congestion, only a sufficient proportion of customers need to be price sensitive within each time-scale that the price can be altered (the rest just contribute more to revenues, reducing the pricing in the long term). This is discussed later. The second counter to this argument is that the introduction of variable pricing where there are third party payers will in itself have implications on how third parties contract to offer payment. For instance, as well as customers running agents that embody their own commercial goals, if someone else is paying, they may do so on condition that their goals (e.g. the company's goals) are also reflected by the agent.

2.2.4 Charging and transmission independence justification

It is wise not to design any differential service for wide area deployment without considering how to charge for it efficiently. Any differential service will need some way to police access to it and, as already discussed, charging is inherently more scalable than policy-based control.

On the other hand, it is unwise to design a charging system for differential services that is tailored to one approach to QoS. From the experience of operational support systems (OSS) for current telecommunications business, development and operations costs are each roughly 50% of total OSS costs [Cochrane98, Readhead98]. If billing systems are focussed on specifically, development drops to about 24% of the total [Cochrane98], but it would still be foolhardy to design a "stove-pipe" solution for each mechanism. This is not to say that the architecture shouldn't include scope for modules specialised to one mechanism (which it does). It is merely to say that such specialised modules should sit in a generic framework for all such modules.

While development is a large part of operational support costs, operational support is a large element of the total cost of running a telecommunications business. The most endearing aspect of the connectionless Internet model is that it could theoretically serve all communication needs with one network to manage [Crowcroft1n?]. Thus it is imperative that there is also only one charging system to manage. If we look to existing telecommunications operators for estimates of the cost of charging, definitions, secrecy and accounting practices make accuracy difficult. However, 39 of BT's 123 billing systems (probably the major ones) cost £121M p.a. (1% of company costs) even without

42

23 billing systems (probably the major ones) cost £121M p.a. (1% of company costs) even without taking the operational staff into account [Cochrane98]. Varian quotes billing and accounting as about 6% [Bailey95] of total costs of an unnamed telecommunications company. BT is believed to be a world-leader in terms of efficiency of its operational support, so billing being 3-5% of BT costs doesn't seem an unreasonable extrapolation of these figures. However, as well as billing, the proposed architecture covers payment handling as well as a large part of network management. 64 of BT's 393 systems covering all these functions cost £179M pa. It is difficult to estimate what proportion of these systems could be amalgamated to support multiple network technologies, but it is possible that running a single network could save about 1-3% of company costs compared to running two, as an extremely rough guess.

The separation between charging and delivered service measurement is invariably absent from the literature. Both sending and receiving applications will need to state what they require of their platform in terms of QoS, multicast or whatever, but don't need or want to consider charging and pricing, which should not be the domain of middleware. Charging is however, the domain of either the customer or the provider themselves or their respective "business rule agents".

This flexibility is particularly important given the turmoil over QoS on the Internet Architecture Board (IAB). RSVP [Zhang93], which was first proposed in 1987, became central to the IETF's integrated services architecture (ISA) [RFC1633], published as a request for comments (RFC) in 1994. As the first (cisco) implementations appeared through 1997 it became clear that the demands of RSVP appeared to conflict heavily with normal router operation, such that an applicability statement had to be made [RFC2208] limiting RSVP to small-scale scope for the foreseeable future. Many architects are claiming in retrospect that RSVP was obviously unscalable and are adopting the new differential services (diff-serv [Nichols98]) architecture with perhaps unseemly haste. Various other QoS proposals are all being considered more plausible to fill the remaining vacuum. Given the limited implementation experience with diff-serv and consequent lack of knowledge of how routers will behave with the new load on them, there is a danger of another RSVP debacle looming, hence the need for the charging system to keep an open mind on what will be charged for.

This worry over undue haste is not just a general note of caution. It will be instructive to digress briefly at this point to provide a critique of one aspect of the differential services architecture specifically related to the separation of charging and transmission. Diff-serv includes a per-packet bit to signal whether each packet is within a usage profile contract or not. This embeds a particular contractual model into the transmission system, which even RSVP didn't do. With a more general contractual model such as that being proposed in this paper, this "in/out" bit would become redundant except where the model of a pre-agreed usage profile was used. Perhaps because of the undue haste, this model appears to have been accepted with very little explicit discussion of whether it is feasible for applications to know their needs in advance or whether it is feasible for a network to combine all these predictions into a correctly dimensioned network over time.

Returning to justifying our approach, the greatest flexibility required of the charging architecture is to be able to stretch between charging for session signalling and for packet-based signalling, e.g. measuring the size of packets with type of service (TOS) bits set. The pay-and-display model is theoretically very capable of doing per-packet charging which has otherwise been consistently scoffed at, although it may well be necessary for some aspects of diff-serv. Per data packet charging is certainly the most general approach, though it is best to try to avoid it where possible.

Other justifications for separating charging from transmission have already been mentioned, in particular the need to allow different ISPs to offer differentiated tariff structures for the same services. The question as to whether over-dimensioning or usage-based charging is cheaper can be decided through competition between ISPs in this way.

2.2.5 Business model independence justification

It would seem as though being able to encapsulate the business of an ISP as a component of a higher level service needs no justification. Placing the network business model interface above the connectionless network layer but below the transport layer may seem obvious, but it does make many higher level business models rather inefficient. Often, the measurement of network service consumption will be at the "wrong" end compared to who is paying for it. Were the boundary above the session layer, charging the session initiator would have a much greater chance of hitting the "correct" end - the party paying. However, this is not how the world is - there is a very definite

business boundary between providers of network and the transport layers for very sound technical reasons embedded in the end-to-end model of the Internet architecture, so there must be a business interface here.

It is likely that, because this is an inherent problem (the situation would be no better without pay-and-display for instance) application business models where every party pays their own edge-network provider directly will predominate over those where one party pays for the network service of others. It is perfectly possible for anyone to pay for anyone else under the architecture, but this is clearly not without overhead. Further justification for the component-based approach has already been given.

3. Architecture

[See Part II]

Part I is intended to be readable without Part II. The Architecture section is written in terms of the components that any implementation would require, rather than the approach used so far, which relies on general description, general principles, underlying concepts and rationale. As well as describing any implementation, this section forms the basis of the design of our own implementation.

4. Related work

This paper has ranged widely so related work is correspondingly across a wide scope, covering:

- network economics & pricing
- network accounting & measurement
- network quality of service
- network access and admission control
- software and information announcement / deployment technologies
- business models and components
- security
- on-line payment technologies

4.1 Network economics & pricing

There are a number of reviews of the network economics literature, including our own [Reeder98] where we examine about twenty of the more influential Internet-related papers. A succinct but useful summary of the approaches from the various constituencies with a stake in the Internet economics debate was written by Bailey following the MIT Internet Economics Workshop in Mar 1995 [BaileyMcK95]. Sarker provides a critique of two leading pricing proposals pointing to problems he proposes (without much proof) would best be solved by regulation [Sarkar95]. Bolot [Bolot94] gives another useful review.

MacKie-Mason and Varian have contributed a number of useful papers explaining network economics such as [MacKieVar95], which makes the simplifying assumption of a single application environment and [Varian96]. Lehr & Weiss analyse the economics of congestion charging across multiple network providers [Lehr96]. Shenker's contributions, which underlie the current paper, have already been discussed at length; that on edge pricing [Shenker96] (after Van Jacobsen [Jacobsen95]) and that on utility curves for applications with various levels of elasticity [Shenker95] that was criticised earlier. Incidentally, Shenker [Shenker95] identified two issues major with usage-based charging - how to charge for multicast and the difficulty of models where the receiver is required to pay. These issues evaporate with the model discussed in our architecture, where payment is from the edges to the middle by both senders and receivers. Sharkey concentrates on what is required to minimise overall cost [Sharkey95].

As well as explaining Internet economics, MacKie-Mason and Varian proposed a specific pricing scheme based on a Vickrey second price auction in their 1992 paper introducing the idea of smart markets [MacKieVar92] that they expanded in 1993 [MacKieVar93]. Unfortunately, although they propose the end-system setting the precedence field in each IP packet, because the price actually charged in such an auction depends on the cut-off level of the router at the time, further messaging is needed if one assumes the user wants or needs price feedback. Thus to manage such an auction

44
needed if one assumes the user wants or needs price feedback. Thus to charge such an auction would require more message round trips in order to implement a full system let alone payment and settlement, so it is not appropriate for per-packet charging. Around the same time, Bohn *et al* [Bohn93] proposed an essentially voluntary system, which uses the precedence field in the IP header, packets with higher precedence being served first. They suggested quotas could be paid for, but essentially left work on policing usage of the service levels for further work. Cocchi *et al* [Cocchi93] successfully experimented with optimising application performance by simulating multiple service classes differentiated by price. No consideration was made of the practicalities of accounting. Other suggestions for specific mechanisms are Crowcroft's proposal for hosts to use multiple network addresses for various service levels [Crowcroft96] and Clark's use of an in/out flag to signal whether a user is within their expected capacity allocation [Clark95] that was developed into an Internet draft [Clark97]. This aspect has already been strongly criticised on the basis it breaks a key architectural separation between charging and transmission, but it has made its way into current IETF differentiated services proposals (diff-serv). Clark's work also suffers from an implicit assumption that a user's requirements must be expressed as a single maximum requirement, rather than being able to have different requirements at different times or per flow (which may be simultaneous) or even per packet. Diff-serv assumes that serendipity can't be given guarantees - that a customer must pre-plan to get guarantees. This is an unfortunately restrictive business model. Already ISPs can and are making a business out of giving statistical guarantees without adding up everyone's planned capacity. Also diff-serv is discussed primarily in terms of dropping packets that are outside the contracted user profile, not in terms of delaying packets (although this is mentioned as an after-thought). This implies the architecture doesn't have the need for differentiated service for interactive real-time applications as its primary focus, which seems strange. Diff-serv requires a pair of packet conditioning elements at every customer provider boundary, all having to be managed. Worse, all these conditioners have to block the packet path while looking up the relevant service profile and deciding whether to clear the bit indicating the packet is in profile. Yet worse still, the service profile databases are likely to contain profiles for very large numbers of customers, making look-ups slow. Also, diff-serv fails to avoid arbitrage while no mechanism exist for service profiles to change dynamically over time or in parallel for simultaneous applications being used by the same customer. Clark also states [Clark95] that he believes differential pricing will lead to hogging, which implies he isn't considering prices that can be varied whenever necessary. Kelly *et al* analyses networks for use by elastic applications and concludes that proportional fairness requirements (such as those supplied by the TCP mechanism) lead naturally to proportionally fair pricing implementations [Kelly98].

4.2 Network accounting & measurement

Ruth provides a succinct but comprehensive and informative review of Internet accounting initiatives [Ruth97]. Basic analysis of cost allocation methods can be found in Peyton-Young [Peyton85]. Herzog analyses how to allocate costs within multicast trees in [Herzog95]. However, it is not at all clear that costs within a network will be exposed externally. That is, due to similar arguments as those leading to "the death of distance" [Cairncross97], it is often cheaper not to account for true costs even though one would expect prices to track these in a truly competitive market [Clark95]. For instance, consideration of what price one would set in reality for the first entrant to a multicast is beyond the scope of Herzog's analysis, despite the impressive proofs he provides to all his assertions.

Therefore, most work tries to find pragmatic factors to measure. Kelly's work identifies two parameters to measure, particularly for price control back-pressure, but also for cost allocation on bursty connections [Kelly97], but schemes that require identification of connections such as Kelly's can be costly to implement on connectionless networks.

However, where currently implemented solutions are concerned, all approaches do indeed involve the identification of flows. Edell *et al* [Edell95] reports on their implementation of a system for the Berkeley campus (40,000 users) that accounts for number of bytes transmitted on a per flow basis. The main problem they solve is that of mapping between flows and customers on multi-user machines in a large user community. A request for the user (or an agent) to authorise charging (just whether it is allowed, not how much) is triggered by the detection of a new TCP connection. The paper concentrates on a single approach rather than exploring alternatives that might avoid the network having to handle transport layer concepts like port numbers. For example, one could map the network address of a host to a primary customer account then delegate charging other users to this account. This could operate through software on the host, supplied by the network provider such that the subcontracting was invisible to all users including the primary account - another example of how the most efficient technical arrangement doesn't have to restrict the commercial arrangements.

45

Despite this, the paper offers a very useful implementation experience for a billing system. MultTCP [Browcroft98] sells the right to break the TCP window increase/decrease rules. However, again because it is dealing with a transport layer concept, it is handicapped. It relies on a proxy host trusted by the link provider to police the service and would not be applicable for a network-wide solution. Also, such transport layer solutions would have scaling problems dealing with UDP flows, which may be long-lived or just single packets (e.g. DNS queries).

To a large extent, accounting for straight bandwidth on a usage basis has been a popular requirement around the Pacific Rim where a small number of expensive shared links supply a large number of fast developing countries. Brownlee's account of the New Zealand experience [Brownlee94] has already been cited. Requirements such as these have led to the creation of an IETF working group on real-time flow measurement (RTFM). Background to the assumptions being made in this group is in request for comments (RFC) 1272 [RFC1272], while their architecture is defined in RFC2063 [RFC2063]. Work in progress (including the architecture) and other related RFCs defining the management information base (MIB) and a simple ruleset language for programming meters are referred to in the group's charter [IETF_rtfm]. The working group is deliberately short-term and doesn't address any more than management control of straightforward bandwidth measurement for assignment against flows, not covering QoS, multicast or sampling etc. It has an unwritten assumption that network providers' measurements will be trusted. It also has to make the dangerous assumption that denial of service will not be a problem, which is a pragmatic assertion for end-users, but less safe for elements of a charging system. Further, the discussion of meter placement [RFC1272] has no assessment of the issue as it applies to a network, merely treating the issue as it applies to a single path through the network. However, despite these shortcomings, as a pure measurement element, RTFM can be fitted into a wider charging architecture without fundamental changes, except for adding sampling and running it over SNMP over TCP rather than UDP transport.

Where QoS is concerned, the most mature mechanism is RSVP, and as such work has started on expressing admission policy for a reservation one aspect of which might be usage-based charging in the RSVP admission policy (RAP) working group. There is no completed work as yet, but all work in progress, including the framework, may be accessed through the group's charter [IETF_rap]. Our own QoS work has involved implementing accounting for RSVP [TasselBri97].

Bulk metrics (of use for capacity planning or price setting for instance) are being standardised in the IP Performance Metrics (IPPM) working group of the IETF. To date, all work is still in progress, including the framework description, and can all be found referenced in the group's charter [IETF_ippm].

The discussions over the international accounting rate systems on the PSTN point to some of the pitfalls experienced in using too simple a model. The PSTN international accounting rate system is described in less formal terms than the previously cited D.150 specification [ITU_D.150] in a useful ITU publication [ITU96], although its characterisation of the Internet model was earlier exposed as flawed. Background to the pressure to change this system can be found on the ITU's Web site [ITU_RIARS].

4.3 Network quality of service

There is a huge body of literature on network QoS but it is outside the scope of this paper to review it here. One of the primary goals of a charging system is that it must be able to be applied to measure and charge for any foreseeable mechanism, therefore this work must attempt to treat QoS mechanisms in as general a sense as possible. Consideration of the application of the architecture to specific mechanisms is covered in the section on the Price Controlled QoS architecture. Implementation and testing of the charging architecture naturally relies on APIs to existing QoS mechanisms which currently boils down to the Integrated Services Architecture (ISA), but with Differentiated Services (diffserv) likely soon. ISA is based on RSVP which is covered by two IETF working groups: Integrated Services (intserv) [IETF_intserv], and RSVP [IETF_rsvp] while the differentiated services working group (diffserv) [IETF_diffserv] covers per packet QoS signalling.

4.4 Network access and admission control

There is relatively little work on access control to classes of service *within* the IP service layer. The Remote Authentication Dial-In User Service (RADIUS) [IETF_radius] only deals with access to the whole IP service although it can account on a finer granularity (note it is targeted at dial-up only).

46
hole IP service although it can't count on a finer granularity (note it is targeted at dial-up only). Radius accounting is very localised and DIAMETER [calhoun98] has been proposed to improve this situation [Tsirtsis98]. However, as has already been stated, this doesn't cater for access control to classes of IP service. The only work in this area is that on local policy modules (LPMs) for RSVP already mentioned under RSVP admission policy (RAP) [IETF rap].

4.5 Software and information announcement / deployment technologies

The architecture described in this paper requires reliable software deployment capabilities. Much of the software required is system software that is difficult to implement in a platform independent way such as with Java [Kramer96] or in a platform dependent but non-portable way such as with Active-X [ActiveX]. Indeed many of the solutions are relatively proprietary. The approach taken by our own work has been to write system specific software where necessary but provide Java wrappers to it. Then most software (e.g. algorithm) updates can be deployed in Java.

The best reference on this field is a comparative review of technologies from the ANSA consortium [Faupe197]. This covers

- the Java applet deployment model
- the Object Management Group's (OMG's) Object Management Architecture (OMA) (which doesn't provide a deployment capability *per se*)
- Oracle's Network Computing Architecture (NCA)
- IBM's Component Broker
- Marimba Castanet and other Web "push" systems
- Mobile agent architectures such as Voyager
- ANSA Flexinet

The following deployment technologies are also mentioned but not analysed in depth:

- Lotus Notes
- Microsoft Open Software Deployment (OSD) and Content Definition Format (CDF)
- Java Beans and Java Electronic Commerce Framework (JECF)

Part of the deployment problem is encountered prior to install time - at build time. The problem is how to add charging components to existing applications without too much complexity for the application programmer (or later by a system integrator). We explored the use of the meta-object protocol to add non-functional aspects like to QoS and charging to functional aspects like unreliable data channels in our paper on QoS [TasselBri97].

As well as deployment of software (e.g. tariff algorithms, electronic contracts), the tariff distribution architecture requires simple data announcement mechanisms too (e.g. price coefficients for tariff algorithms or addresses of new channels). Handley's session announcement protocol (SAP) [HandleySAP??] forms the basis of these mechanisms. SAP uses Handley's session description protocol (SDP) [HandleySDP??] as the format to describe sessions. However, the work by Evans on a more generic data distribution mechanism called reliable scalable, secure distribution (RSSD) [Evans98??] and work by Eyles on a more generic description format called new-fangled announcement framework (NFAF) [Eyles98??] is necessary in order to stretch Handley's work so that sessions can be abstracted to encompass tariffs. Further, it has already been described how it might be necessary to add non-repudiable delivery timing if these announcements become very frequent. This was solved using the technique we describe in [Fairman98].

4.6 Business models and components

The section on Flexible Business Models has already given some references on Internet business models and component architectures. Our literature review discusses some more [Reeder98].

The primary technical forum covering business objects is the Business Object Design and Implementation Workshop, held annually [OMG bow97].

The mission of the Object Management Group includes the specification of interfaces to common business objects [OMG borfp96, Shelton97]. Freestone & Owen [Freestone97] describe the technical approach necessary to create a market in business services and also introduce the possibility of fully

47

dynamic business systems all reacting with each other in real time in their section on new opportunities. Sundaresan *et al* [Sundaresan97] give a detailed analysis of the market for sale of BT's systems capabilities while Readhead & Cochrane [Readhead97] give a briefer review including discussion of various attempts to outsource the production of operational support systems. Both draw fairly negative conclusions which reflects the problem with a legacy of solutions already in operation that were never built with such ideas in mind.

4.7 Security

This paper doesn't go into depth on security mechanisms, so it would be inappropriate to refer to literature on this subject. However, the pay-and-display approach is an classic example of the trend away from formally secure systems to pragmatic security with features such as imperfect detection and adaptability akin to the approaches discussed in Somayaj *et al* [Somavaji97].

Our scalable technique for adding non-repudiation to a multicast has already been mentioned [Fairman98]. This technique can also be used to ensure key management for encrypted multicast tariffs only affects the relevant customers. More traditional techniques, which don't require smart cards but which are all less successful at confining the effects of multicast key management to single users are reviewed by Blunden [Blunden98?]

4.8 On-line payment technologies

An introduction to electronic payment systems is provided by Putland *et al* [Putland97], particularly focussing on pre-paid systems as well as fine payment granularity. Underlying this introduction was work that resulted in a useful classification of pre-paid micropayment technologies [Hill96].

/ . References

48

- [{3b?}] {3 bits for charging - via Crowcroft}
- [{Aboba?}] {Discussion on flaws in use of UDP for measurement for settlement, Bernard Aboba}
- [{ACL_perf?}] {Firewall 10% performance drop per ACL, via Dave Douglas}
- [ActiveX] "Microsoft Internet Explorer 4.0 White Paper; Complete Administration", Microsoft, Jul 1997, <URL:<http://www.microsoft.com/ie/press/whitepaper/iwhite/white016.htm>>
- [Axelrod84] R Axelrod, "The Evolution of Co-operation", Harper-Collins, 1984
- [Bailey95] J. Bailey, S. Gillett, D. Gingold, B. Leida, D. Melcher, J. Reagle, J. Roh, R. Rothstein, and G. Seale, "Internet Economics Workshop Notes," Research Program on Communications Policy, MIT, Mar 1995. <URL:<http://www.press.umich.edu/jep/works/BailWNotes.html>>
- [BaileyMcK95] Joseph P Bailey and Les W McKnight, (MIT), "Internet Economics: What Happens When Constituencies Collide", Apr 1995, <URL:<http://aleph.pangea.org/HMP/PAPER/123/html/paper.html>>
- [Barns89] Barns, W, "Defense Data Network Usage Accounting Enhancement Approaches," The MITRE Corporation, 1989.
- [BlueSquirrel] WebWhacker and other Blue Squirrel Web robot products, <URL:<http://www.bluesquirrel.com/products.html>>
- [Bohn93] Roger Bohn(Uni of CA-San Diego), Hans-Werner Braun & Kimberly C Claffy (San Diego Supercomputer Center) and Stephen Wolff (NSF DNCRI), "Mitigating the coming Internet crunch: multiple service levels via Precedence", Technical Report, University of California-San Diego, San Diego Supercomputer Center and NSF, Nov 1993, <URL:<http://www.nlanr.net/Papers/mcic.html>>
- [Bolot94] Jean-Chrysostome Bolot (INRIA), "Cost-quality tradeoffs in the Internet", Computer Networks and ISDN Systems, 1994, <URL:<ftp://www.inria.fr/rodeo/bolot/94.Tradeoffs.ps.gz>>
- [Braden97] B. Braden, D. Clark, J. Crowcroft, B. Davie, S. Deering, D. Estrin, S. Floyd, V. Jacobson, G. Minshall, C. Partridge, L. Peterson, K. Ramakrishnan, S. Shenker, J. Wroclawski, L. Zhang, "Recommendations on Queue Management and Congestion Avoidance in the Internet", IRTF Internet draft, Mar 1997 <URL:<draft-irtf-e2e-queue-mgt-00.txt>>
- [Briscoe96?] Bob Briscoe (BT), "Subject: Re: FW: metering", Review of RTFM architecture (e-mail to private distribution list, BT access only), 18 Apr 1996
- [Briscoe97] Bob Briscoe, QoS-based charging Task plan, Dec 97, work in progress, <URL:<http://www.jungle.bt.co.uk/projects/lsma/charging/admin/9899prop.html>> (BT access only).
- [Brownlee94] Nevil J. Brownlee (Uni. of Auckland), "New Zealand experiences with network traffic charging", usage-based access charging, ConneXions Volume 8, No. 12, Dec 1994, <URL:<http://www.press.umich.edu/jep/works/BrownNewZe.html>>
- [Cairncross97] Frances Cairncross, "The Death of Distance : How the Communications Revolution

"I Change Our Lives", Harvard Business School Press, ISBN 0875848060, Oct 1997

[Calhoun98] Pat R. Calhoun (Sun), Allan C. Rubens (Ascend), "DIAMETER: Base Protocol", work in progress: Internet Draft, Mar 1998, <URL:[draft-calhoun-diameter-02.txt](#)> (also see [Internet Drafts index](#) for many other DIAMETER drafts)

[Carpenter96] Brian Carpenter, "Metrics for Internet Settlements", EXPIRED Internet Draft, May 1996, but the main information from this draft is recorded by Claffy at
<URL:<http://www.nlanr.net/Papers/telegeog96.html>>

[Cisco] Cisco Systems, "Queuing", informal white paper,
<URL:<http://mmgwp.cisco.co.ip/technologvdoc/queue.html>>

[Clark95] David D Clark (MIT), A Model for Cost Allocation and Pricing in the Internet, presented at MIT Workshop on Internet Economics, Mar 1995,
<URL:<http://www.press.umich.edu/jep/works/ClarkModel.html>>

[Clark97] David D Clark and John Wroclawski (MIT), "An Approach to Service Allocation in the Internet", EXPIRED Internet Draft, Jul 1997,
<URL:<http://diffserv.lcs.mit.edu/Drafts/draft-clark-diff-svc-alloc-00.txt>>

[Cocchi93] Ron Cocchi (Hughes/USC), Scott Shenker(Xerox PARC), Deborah Estrin (USC) and Lixia Zhang (Xerox PARC), "Pricing in Computer Networks: Motivation, Formulation, and Example", in: IEEE/ACM Transactions on Networking, vol. 1 (6), pp. 614-627, Dec 1993
<URL:<http://www.wiwi.hu-berlin.de/~brupp/pricing/PricingInComputerNetworks.ps.gz>>

[{CrowcroftIn?}] {Advantage of one network to manage - via Crowcroft}

[Crowcroft96Kn] Jon Crowcroft, "QoS - Plenty of \$\$\$ but just 1 Bit", Keynote speech presented at Protocols for High Speed Networks '96, Oct 1996,
<URL:<http://www.cs.ucl.ac.uk/staff/jon/hipparch/dollarbit/>> based on HIPPARCH discussions between Christophe Diot, Jean-Chrysostome Bolot, Jon Crowcroft & David Clark

[Crowcroft96] Crowcroft, Jon, "Pricing the Internet", in IEE Colloquium on Charging for ATM (ref. no. 96/222) 12 Nov 1996, pp1/1-4.

[Crowcroft98] Jon Crowcroft & Philippe Occhislin (UCL), "Differentiated End to End Internet Services using a Weighted Proportional Fair Sharing TCP", Apr 1998,
<URL:<ftp://cs.ucl.ac.uk/darpa/multtcp.ps>> or <URL:
<http://www.cs.ucl.ac.uk/staff/J.Crowcroft/hipparch/pricing.html>>

[Cochrane98] James Cochrane, Alan Readhead, Pete Barnsley, Maff Holladay, Michael Lyons. (BT). "OSS Costs and Revenues: An Indicative Study" Apr 1998
<URL:http://trafalgardrake.bt.co.uk/pete_grp/deliverable/1997/kev/kd1issue1.html> (BT access only)

[Danielsen95] K. Danielsen (Molde College, No) and M. Weiss (Uni. of Pittsburgh), "User Control Modes and IP Allocation," Journal of Electronic Publication, University of Michigan Press, <URL:<http://www.press.umich.edu/jep/works/DanieContr.html>>, 1995.

[Edell95] Richard Edell (Uni of CA-Berkeley), Nick McKeown (Stanford), and Pravin Varaiya (Uni of CA-Berkeley), "Billing Users and Pricing for TCP." IEEE JSAC Special Issue on Advances in the Fundamentals of Networking, Sep 1995, <URL:<http://tin-v-tera.stanford.edu/~nickm/papers.html>>

[Fairman98] Ian Fairman & Bob Briscoe (BT), "Non-repudiation and Multicast Key Management Using Securely Encapsulated Key Generators", work in progress
<URL:<http://www.jungle.bt.co.uk/projects/lsma/charging/content/SecureKeyGenerators-0.2.html>>
(BT access only)

[Faupel97] Matthew Faupel (APM), "Java Distribution and Deployment", ANSA report APM.2038, Oct 1997 <URL:<http://www.ansa.co.uk/ansasponsors/phase3-doc-root/sponsors/APM.2038.02.ps.gz>> (ANSA sponsor access only)

ANSA sponsor access only)

50

[Finlayson98] Ross Finlayson, Discussion at Third Reliable multicast research group meeting, Orlando, FL, and following on the mailing list, Feb 1998 <URL:<http://www.east.isi.edu/rmv/>>

[{Floyd?}] {Absorbing bursts with one network - via Crowcroft or Floyd}

[Floyd93] Sally Floyd and Van Jacobson, (LBL) "Random Early Detection Gateways for Congestion Avoidance", IEEE/ACM Transactions on Networking, Vol.1, No. 4, pp. 397-413. Aug 1993, <URL:<ftp://ftp.ee.lbl.gov/papers/early.ps.gz>>

[Floyd94] Sally Floyd (LBNL), "TCP and Explicit Congestion Notification", ACM Computer Communication Review, V. 24 N. 5, Oct 1994, p. 10-23. [This issue of CCR incorrectly has "1995" on the cover instead of "1994".] <URL:<http://www.nrg.ee.lbl.gov/nrg-papers.html>>

[FollowMe] "FollowMe", ANSA/ESPRIT research project, <URL:<http://www.ansa.co.uk/Research/FollowMe.htm>>

[Freestone97] D Freestone and M Owen, (BT), "Operational support systems for the future", BT Technology Journal, Vol.15, no.1, p172, Jan 1997. <URL:<http://www.labs.bt.com/library/bttj/v15n1/16.htm>> (abstract & order details only) or <URL:http://sss.info.bt.co.uk/BOOKS_AND_JOURNALS/BTTJ/v15n1/16.htm> (BT access only)

[Herzog95] Shai Herzog (IBM), Scott Shenker (Xerox PARC), Deborah Estrin (USC/ISI), "Sharing the cost of Multicast Trees: An Axiomatic Analysis", in Proceedings of ACM/SIGCOMM '95, Cambridge, MA, Aug. 1995, <URL:<http://www.research.ibm.com/people/h/herzog/sigton.html>>

[Hill96] Jake Hill & Dimitrios Tsapakidis, (BT), "Prepaid Micropayments, A review and generalisation of some practical prepaid micropayment protocols" (draft), 1996 <URL:<http://www.alien.bt.co.uk/~dimitris/mp.doc>> (BT access only)

[Holladay98] Unpublished, incomplete background data to [Cochrane98]

[Huitema95] Christian Huitema, "Routing in the Internet", Prentice Hall, ISBN: 0131321927, Apr 1995

[IEEE802.1p] IEEE, "Draft Standard for Traffic Class and Dynamic Multicast Filtering Services in Bridged Local Area Networks (Draft Supplement to 802.1D)", IEEE802.1p, 1997

[IETF_diffserv] Differentiated Services (diffserv) IETF working group charter <URL:<http://www.ietf.org/html.charters/diffserv-charter.html>>

[IETF_intserv] Integrated Services (intserv) IETF working group charter <URL:<http://www.ietf.org/html.charters/intserv-charter.html>>

[IETF_ippm] IP Performance metrics (IPPM) IETF working group charter <URL:<http://www.ietf.org/html.charters/ippm-charter.html>>

[IETF_radius] Remote Authentication Dial-In User Service (radius) IETF working group charter <URL:<http://www.ietf.org/html.charters/radius-charter.html>>

[IETF_rap] RSVP Admission Policy (RAP) IETF working group charter <URL:<http://www.ietf.org/html.charters/rap-charter.html>>

[IETF_rsvp] Resource Reservation Setup Protocol (rsvp) IETF working group charter <URL:<http://www.ietf.org/html.charters/rsvp-charter.html>>

[IETF_rtfm] Realtime Traffic Flow Measurement (rtfm) IETF working group charter <URL:<http://www.ietf.org/html.charters/rtfm-charter.html>>

[IGMPv3] Brad Cain (Bay), Steve Deering (cisco), Ajit Thyagarajan (Torrent), "Internet Group Management Protocol, Version 3", Work in progress: IETF Internet Draft, Nov 1997 (expires May

8), <URL:draft-ietf-idn-imp-v3-00.txt>

51

[IPv6] S. Deering, (cisco), R. Hinden (Ipsilon), "Internet Protocol, Version 6 (IPv6) Specification". work in progress: IETF Internet draft, Nov 1997, <URL:draft-ietf-ipngwg-ipv6-spec-v2-01.txt> or latest draft through IPng working group charter, <URL:<http://www.ietf.org/html.charters/ipngwg-charter.html>>

[ITU_D.150] ITU-T, "New system for accounting in international telephony", ITU-T Rec. D.150, Series D: General tariff principles - Charging and accounting in the international telephone service; Oct 1996, <URL:<http://www.itu.ch/intset/itu-t/d150/d150.htm>>

[ITU_RIARS] ITU "Reforming the International Accounting Rate System" <URL:<http://www.itu.ch/intset/>>

[ITU96] ITU, "The Direction of Traffic", ITU/Telegraphy Inc, Geneva, 1996 <URL:<http://www.itu.ch/ti/publications/traffic/direct.htm>> in brief chapter 3, Box 3.1 is extracted on-line: "Accounting rates and how they work", <URL:<http://www.itu.ch/intset/whatare/howwork.html>>

[Jacobsen88] Van Jacobsen (LBL), "Congestion avoidance and control", In Proc. ACM SIGCOMM'88, 314-329, <URL:[ftp://ftp.ee.lbl.gov/papers/congavoid.ps.Z](http://ftp.ee.lbl.gov/papers/congavoid.ps.Z)>

[Jacobsen95] Referred to in [Shenker96] as "Van Jacobsen [LBL], private communication, 1995."

[{Kelly?}] {Kelly? Statistical analysis of elastic/inelastic utility curves}

[Kelly97] Frank P Kelly (Cambridge Uni.), "Charging and Accounting for Bursty Connections", in "Internet Economics" (Editors Lee W. McKnight and Joseph P. Bailey) MIT Press, 1997. 253-278. <URL:<http://www.statslab.cam.ac.uk/~frank/charge.html>>

[KellyDisc97] Frank Kelly (Uni. of Cambridge), unpublished discussion at 4th COST237 workshop, Lisboa, Dec 1997

[Kelly98] Frank P Kelly, Aman K. Maulloo and David KH Tan, "Rate control for communication networks: shadow prices; proportional fairness and stability", Journal of the Operational Research Society, 49, 1998, <URL:<http://www.statslab.cam.ac.uk/~frank/rate.html>>

[Kramer96] Douglas Kramer, "The Java™ Platform", Sun White Paper, May 1996, <URL:<http://www.javasoft.com/docs/white/platform/CreditsPage.doc.html>>

[Lehr96] William H Lehr (Columbia Uni.) & Martin BH Weiss (Uni. of Pittsburgh), "The political economy of congestion charges and settlements in packet networks", Telecommunications Policy, 20 (1996), 219-231,

[Lin97] Dong Lin and Robert Morris (Harvard Uni) "Dynamics of Random Early Detection", in Proc. SIGCOMM'97, Computer Communication Review, ACM SIGCOMM, volume 27, number 4, ISSN # 0146-4833, Oct 1997 <URL:<http://www.inria.fr/rodeo/sigcomm97/program.html#ab078>>

[MacKieVar92] Jeffrey K MacKie-Mason and Hal Varian, (UMich), "Some Economics of the Internet", Tenth Michigan Public Utility Conference at Western Michigan University, March 25-27, 1993: <URL:<http://www.sims.berkeley.edu/~hal/people/hal/papers.html>>

[MacKieVar93] Jeffrey K MacKie-Mason and Hal Varian, (UMich), "Pricing the Internet", Public Access to the Internet, JFK School of Government, Apr 1993 <URL:<http://www.sims.berkeley.edu/~hal/people/hal/papers.html>>

[MacKieVar95] Jeffrey K MacKie-Mason and Hal Varian, (UMich), "Pricing Congestible Network Resources", IEEE Journal on Selected Areas in Communications, "Advances in the Fundamentals of Networking", 1995, <URL:<http://www.sims.berkeley.edu/~hal/Papers/pricing-congestible.pdf>>

[Newman96] Peter Newman, Tom Lyon & Greg Minshall, (Ipsilon) "Flow Labelled IP: A

52
Connectionless Approach to ATM Proc. IEEE Infocom'96, vol. 3, pp. 1-60, Mar 1996
<URL:<http://www.ipsilon.com/~pn/papers.html>>

[Nichols98] Kathleen Nichols (Bay), Steven Blake (IBM), (eds), "Differentiated Services Operational Model and Definitions", work in progress, IETF Internet Draft, Feb 1998
<URL:<http://www.ietf.org/internet-drafts/draft-nichols-dsopdef-00.txt>> (see [IETF_diffserv] for latest version)

[Odlyzko97] Andrew Odlyzko (AT&T Research), "A modest proposal for preventing Internet congestion", Sep 1997, <URL:<http://www.research.att.com/~amo/doc/recent.html>>

[OMG_borfp96] OMG request for proposals (RFP) for Common Business Object Specifications
<URL:<http://www.omg.org/pr96/bomrfp.htm>>

[OMG_bow97] "OOPSLA'97 Business Object Workshop" Forthcoming, Proceedings, pub. Springer-Verlag <URL:<http://www.tiac.net/users/jsuth/oopsla97/index.html>>

[Peyton85] H. Peyton-Young, editor, "Cost Allocation Methods, Principles, Applications. pub. Elseviers, 1985

[Putland97] Paul Putland, Jake Hill, and Dimitrios Tsapakidis, "Electronic payment systems", BT Technology Journal, Vol.15, no.2, p32, Apr 1997.
<URL:<http://www.labs.bt.com/library/bttj/v15n2/03.htm>> (abstract & order details only)
<URL:http://sss.info.bt.co.uk/BOOKS_AND_JOURNALS/BTTJ/v15n2/03.htm> (BT access only)

[Ramakrishnan97] K. K. Ramakrishnan (AT&T Labs Research) and Sally Floyd (LBNL), "A Proposal to Add Explicit Congestion Notification (ECN) to IPv6 and to TCP", Internet Draft <draft-kksjf-ecn-00.txt>, Nov 1997, <URL:draft-kksjf-ecn-00.txt>

[Readhead97] Alan Readhead and James Cochrane, "OSS Revenue Opportunities, What new revenues could OSS bring to BT?", Oct 97 (BT access only)

[Readhead98] {OSS cost mostly development, Alan Readhead}

[Reeder98] Tony Reeder & James Cochrane (BT), "A Review of Internet Charging Models", Mar 1998,
<URL:<http://www.jungle.bt.co.uk/projects/lisma/charging/qos-based/busmodel/review/review~1.html>> (BT access only)

[RFC768] Jon Postel (USC/ISI), "User Datagram Protocol", STD 6, IETF RFC 768, Aug 1980,
<URL:rfc768.txt>

[RFC791] Jon Postel (ed), "Internet Protocol", IETF RFC 791, Sep 1981, <URL:rfc791.txt>

[RFC793] Jon Postel (ed), "Transmission Control Protocol", STD 7, IETF RFC 793, Sep 1981,
<URL:rfc793.txt>

[RFC1272] C. Mills, D. Hirsh, G. Ruth, "Internet Accounting: Background" IETF RFC 1272, Nov 1991 <URL:[ftp://ftp.isi.edu/in-notes/rfc1272.txt](http://ftp.isi.edu/in-notes/rfc1272.txt)>

[RFC1633] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet architecture: an overview", IETF RFC 1633, Jun 1994. <URL:<http://www.isi.edu/div7/rsvp/pub.html>>

[RFC1905] J. Case, (SNMP Research), K. McCloghrie (cisco), M. Rose (Dover Beach Consulting) & S. Waldbusser (International Network Services), "Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)", IETF RFC 1905, Jan 1996, <URL:rfc1905.txt>

[RFC1906] J. Case, (SNMP Research), K. McCloghrie (cisco), M. Rose (Dover Beach Consulting) & S. Waldbusser (International Network Services), "Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2)", IETF RFC 1906, Jan 1996, <URL:rfc1906.txt>

[RFC2001] W. Stevens (NEC), "TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms" IETF RFC2001, Jan 1997 <URL:[rfc2001.txt](http://www.ietf.org/rfc/rfc2001.txt)>

[RFC2005] J. Solomon (Motorola), "Applicability Statement for IP Mobility Support", IETF RFC 2005, Oct 1996, <URL:[rfc2005.txt](http://www.ietf.org/rfc/rfc2005.txt)>

[RFC2063] N. Brownlee, C. Mills, G. Ruth, "Traffic Flow Measurement: Architecture" IETF RFC 2063, Jan 1997 <URL:[rfc2063.txt](http://www.ietf.org/rfc/rfc2063.txt)>

[RFC2208] F. Baker, B. Braden, S. Bradner, A. Mankin, M. O'Dell, A. Romanow, A. Weinrib, L. Zhang, "Resource ReSerVation Protocol (RSVP) Version 1 Applicability Statement Some Guidelines on Deployment", IETF RFC 2208, Jan 1997, <URL:[rfc2208.txt](http://www.ietf.org/rfc/rfc2208.txt)>

[RFC2236] B. Fenner (Xerox PARC), "Internet Group Management Protocol, Version 2", IETF RFC2236, Nov 1997, <URL:[rfc2236.txt](http://www.ietf.org/rfc/rfc2236.txt)>

[Ruth97] Gregory R. Ruth (GTE), "Usage Accounting for the Internet", In Proc. INET'97 Kuala Lumpur, <URL:http://www.isoc.org/isoc/whatis/conferences/inet/97/proceedings/F1/F1_1.HTM>

[Sarkar95] Mitrabarun Sarkar, "An Assessment of Pricing Mechanisms for the Internet - A Regulatory Imperative," Journal of Electronic Publication, University of Michigan Press, Mar 1995, <URL:<http://www.press.umich.edu/jep/works/SarkAssess.html>>

[Sharkey95] W W Sharkey, "Network Models in Economics", in MO Ball *et al* (eds), "Handbook of Operations Research and Management Science: Networks", Vol 8, pp713-765, 1995

[Shelton97] Robert Shelton ed., "Common Business Objects", bom/97-12-04 Object Management Group white paper, <ftp://ftp.omg.org/pub/docs/bom/97-12-04.pdf>

[Shenker95] Scott Shenker (Xerox PARC), "Fundamental Design Issues for the Future Internet", IEEE Journal on Selected Areas in Communications, 1995. URL: <<http://www.spp.umich.edu/spp/courses/744/docs/FundamentalDesign-shenker.pdf>>

[Shenker96] Scott Shenker (Xerox PARC), David Clark (MIT), Deborah Estrin (USC/ISI) and Shai Herzog (USC/ISI), 'Pricing in Computer Networks: Reshaping the research agenda', SIGCOMM Computer Communication Review Volume 26, Number 2, Apr 1996, <URL: <http://www.statslab.cam.ac.uk/~frank/PRICE/scott.ps>>

[Skevington97] Peter J Skevington and Tim P Hart (BT), "Trusted third parties in electronic commerce", BT Technology Journal, Vol.15, no.2, p39, Apr 1997, <URL:<http://www.labs.bt.com/library/bttj/v15n2/04.htm>> (abstract & order details only) or <URL:http://sss.info.bt.co.uk/BOOKS_AND_JOURNALS/BTTJ/v15n2/04.htm> (BT access only)

[Somayaji97] A. Somayaji, S. Hofmeyr, and S. Forrest, "Principles of a Computer Immune System" New Security Paradigms Workshop, Sep 1997 <URL:<http://www.cs.unm.edu/~steveah/papers.html>>

[Speakman98] Tony Speakman, Dino Farinacci, Steven Lin, Alex Tweedly, (cisco) "PGM Reliable Transport Protocol Specification", Work in progress: IETF Internet Draft, Jan 1998 (expires Jul 1998), <URL:[draft-speakman-pgm-spec-01.txt](http://www.ietf.org/drafts/speakman/pgm-spec-01.txt)>

[Stoica97] Ion Stoica and Hui Zhang, (CMU), "A Hierarchical Fair Service Curve Algorithm For Link-Sharing, Real-Time and Priority Services", in Proc. SIGCOMM'97, Computer Communication Review, ACM SIGCOMM, volume 27, number 4, ISSN # 0146-4833, Oct 1997 <URL: <http://www.inria.fr/rodeo/sigcomm97/program.html#ab011>>

[Sundaresan97] Hari Sundaresan *et al*, "Selling BT's System Capabilities", Telecommunications Master's Programme 6, Aug 1997 (BT access only)

[Tassel97] Jérôme Tassel (Uni of Kent at Canterbury), "Quality of Service (QoS) adaptation using Reflective Java", MSc dissertation, Sep 1997,

54

URL:http://www.labs.bt.com/people/briscorj/projects/lisma/jt_thesis/thesis_final.htm>

[TasselBri97] Jérôme Tassel, Bob Briscoe, Alan Smith, (BT), "An End to End Price-Based QoS Control Component Using Reflective Java", in Lecture Notes in Computer Science from the 4th COST237 workshop, pub. Springer-Verlag, Dec 1997,
<URL:<http://www.labs.bt.com/people/briscorj/papers.html#QoteS>>

[Tatham98] Martin Tatham (BT), "A Commercial Model for IP Multicast", Feb 1998
<URL:http://gideon.bt-svs.bt.co.uk/martin/multicast/mcast_v1.doc> (BT access only)

[Thomas96] Stephen A. Thomas, "IPng and the TCP/IP Protocols : Implementing the Next Generation Internet", pub Wiley, ISBN: 0471130885, Jan 1996

[Tsirtsis98] George Tsirtsis (BT), "RADIUS - DIAMETER Comparison", March 1998
<URL:<http://www.gideon.bt.co.uk/george/radvsdiam/RADvsDIAM.htm>> (BT access only)

[U-Scan] Optimal Robotics, "U-Scan", self-checkout product,
<URL:<http://www.optimal-robotics.com/index.htm>>

[Varian96] Hal R Varian, "Differential Pricing and Efficiency", f i ® s T - m o n d @ ¥, Vol.1 No.2, Aug 1996, <URL:<http://www.firstmonday.dk/issues/issue2/different/>>

[Xiaoming96] Hao Xiaoming (Nanyang Technological University, Singapore), Huang Yu (Baptist University, Hong Kong) & Kewen Zhang (University of Missouri), "The Internet and Information Control: The Case of China", in Proc. INET'96,
<URL:http://cios.llc.rpi.edu/getfile/Xiaoming_V6N296> (subscription required)

[Zhang93] Lixia Zhang (Xerox PARC), Stephen Deering (Xerox PARC), Deborah Estrin(USC/ISI), Scott Shenker (Xerox PARC) and Daniel Zappala (USC/CSD), "RSVP: A New Resource ReSerVation Protocol", IEEE Network. Sep 1993. <URL:<http://www.isi.edu/div7/rsvp/pub.html>>

[Zull97] Chris Zull (Cutler & Co), "Interconnection Issues in the Multimedia Environment", Interconnection Asia '97, IIR Conferences, Singapore, Apr 1997
<URL:<http://www.cutlerco.com.au/core/content/speeches/Interconnection%20Issues/Interconnection>>

Notes

(i) Examples of potentially dynamically switched network technologies, over which IP may be "tunnelled" include:

- switched multi-megabit data service (SMDS)
- asynchronous transfer mode (ATM) switched virtual circuits (SVCs)

However, when such solutions are used, it is usually only a short-term phenomenon, e.g. to utilise existing investment or to take advantage of a facility yet to be supplied by native IP routers such as certain management interfaces on the routing equipment. Given the concentration of worldwide development effort on native IP solutions, these scenarios are likely to decrease in number, so are relegated to be the subject of further work.

(ii) There is a possibility that bandwidth and latency can be reduced to a single dimension. Latency can be considered as bandwidth delivered within a specific time as opposed to bandwidth delivered over continuous time[Stoica97].

Although reliability is definitely a separate dimension, there is a large sub-set of requirements where QoS dimensions can be simplified by assuming virtually 100% reliability which can be achieved where high QoS classes of service are layered over best-effort. However there are very large classes of requirement where QoS without high reliability is acceptable e.g. with controlled loss transmission of audio and video.

For instance:

- a business could buy Internet service from different providers with differing levels of over-dimensioning and price and re-sell them all as a package to a customer wanting switchable levels of service
- or one business could operate an ATM network and buy basic IP service (but built on IP switches[Newman96]) from another ISP such that packets to or from customers of the former would cut through the ATM network

(iv) Exceptions to flat-fee access charging do exist, such as the experience in New Zealand of usage-based charging[Brownlee94] and that in Italy, which lost out to flat-fee charging through competition[BaileyMcK95].

(v) The costs of packet duplication (multicast) can be covered by flat charges for access bandwidth, but while competition in the provision of multicast service is ineffective (which may be a long time), ISPs will want to charge differentially for the added value of multicast despite this [Tatham98].

(vi) For multicast, it is already not possible to send data to someone unsolicited because the Internet model is receiver initiated. Even for a multi-source multicast, per-source joins have recently been proposed for v3 of the Internet group management protocol (IGMP)[IGMPv3] so that a receiver can selectively cut out certain senders and still receive the rest of the multicast.

For aggregation, the "receiver" (of aggregated reverse multicasts) at the top of the tree will always have set up the tree in the first place. The scheme must be carefully designed to avoid this "receiver" being open to the highest charge any sender can impose on them. This is the case for the only known mature example of aggregation, the reservation protocol (RSVP)[Zhang93]. Here "data receiver" initiated reservation (RESV) messages are aggregated, each router always only passing on the highest QoS specification (RSPEC) but with an upper limit set by the TSPEC in the PATH messages sent continuously in the opposite direction by the "data sender" (or "receiver" of RESV messages) at the top of the tree. Incidentally, the data sender is notified when the aggregated reservation results in a change at the top of the tree, by a INFO_RESV message.

(vii) Packet drop can be for a number of reasons, most of which involve silent drop when queues or buffers overflow. However, the list also includes losses that are reported as errors using ICMP such as when a router has no route for a destination address or the time to live expires, possibly due to a routing loop. It is possible, though by no means certain, that, as QoS mechanisms become standardised, when non-best-effort packets are dropped, an (unreliable) error will have to be returned to the sender.

(viii) Examples of packets that are forwarded until aggregation (reverse multicast) are:

- RSVP[Zhang93] receiver initiated reservation (RESV) messages
- pragmatic general multicast (PGM) [Speakman98] negative acknowledge (NACK) messages or the "lay breadcrumb" messages[Finlayson98] suggested in their place

~~(ix) The position has now changed because the obligation has been removed, so the letter may simply be returned to the sender.~~

(x) The lack of success of anti-spamming measures is primarily because there is not a spectrum of punishments available at sensible cost. An ISP either has to kick off a spammer (losing an otherwise valuable customer) or ignore it. If the infrastructure were in place for making small payments and fines, an ISP could efficiently fine senders of unsolicited information to the point where it became uneconomic for the individual to continue. The ISP would then gain revenue rather than lose it due to the punishment.

(xi) This is not to say that all conferences don't include chargeable information or that conference organisers should never charge, even for organising a valuable gathering.

(xii) To clarify, this is *not* to say that any party participating in a higher level application involving usage-based charging:

- 56
- won't be involved in commercial interactions with the other parties - the application must be paying their ISP directly (someone else may be paying for them, possibly through a clearing system)
 - can't be paying other edge ISPs directly too
 - can't be paying other backbone ISPs directly too

However, what it *is* saying is that it is in an ISP's interest to at least *offer* a "one-stop shop".

(xiii) There is further irony in the difficulty with which many large corporations in the "free-world" have accepted the new-found uncontrolled access to cheap information publication by their employees {ref?}.

(xiv) By the same argument, there is no such thing as a non-real-time application.

The term "real-time" applications, as used in this context should strictly be "soft-real time" and by the same arguments, all applications are real-time to varying degrees. Hard real-time means there are non-negotiable deadlines (e.g. synchronous control systems) whereas soft real-time means the deadlines may be achieved statistically.

(xv) Fortunately for BT's internal e-mail system :-)

(xvi) In the case of latency, there is an absolute limit for any particular network diameter because of the speed of light, so the curve starts to apply to smaller and smaller networks in this case. However, there are no such absolute limits for other dimensions of QoS, so in general this aspect can be ignored.

(xvii) As the pay-as-you-go model shrinks the window during which records need to be kept, "it would be nice" to rely less on stable storage. However, this architecture doesn't go as far as proposing such a small window that it becomes uneconomic (compared to the revenue loss risk) to cater for recovery of transactions in progress after a system crash.

CLAIMS

1. A method of operating a communications network including
distributing a tariff via a communications network to a multiplicity of
5 customer terminals connected to the communications network, and
calculating using the said tariff a charge for use by the customer terminal
of the network to which the tariff applies.
- 10 2. A method according to claim 1, in which the tariff algorithm is distributed to
the multiplicity of customer terminals via the communications network to which
the said tariff applies.
- 15 3. A method according to claim 1 or 2, in which the step of distributing the tariff
includes steps of communicating separately a formula for calculation of network
usage charges , and coefficients for use in the said formula.
4. A method according to any one of the preceding claims, including a further step
20 of distributing to the customer terminals a revised tariff.
5. A method according to claim 4 when dependent on claim 3, in which the step
of distributing a revised tariff comprises communicating revised coefficients for use
in the formula previously distributed to the customer terminals.
- 25 ~~6. A method according to claim 4 or 5, including detecting loading of network
resources and determining a revised tariff in dependence upon the results of the
said step of detecting loading.~~
- 30 7. A method according to claim 6, in which the steps of detecting loading and
determining a revised tariff are carried out automatically by a network management
platform.

8. A method according to anyone of the preceding claims including communicating to a customer terminal data identifying a first predetermined communications channel, and at the customer terminal subsequently monitoring the said communications channel for communications relating to the said tariff.

5

9. A method according to claim 8, including communicating on the said first communications channel data identifying one or more further communications channels, and the customer terminal subsequently monitors in addition the or each further channel.

10

10. A method according to claim 9, including introducing a new communications channel and identifying the said new communications channel on a communications channel previously identified to the customer terminal depending on loading of the said previously identified communications channel.

15

11. A method according to any one of the preceding claims including communicating encrypted tariff data to the customer terminal, and decrypting the said tariff data within a secure module located at the customer terminal.

20

12. A method according to claim 11 including communicating different tariff data on a plurality of different communication channels and providing at a customer terminal a key specific to tariff data on one of the plurality of communication channels.

25

13. A method according to any one of the preceding claims, including operating a plurality of different services on the communications network, communicating different tariffs for different respective services to the multiplicity of customer terminals, and selectively varying a respective tariff depending on an operational condition of the respective service.

30

14. A method of operating a communications network comprising:
operating a plurality of different services on the network;

communicating tariffs for the different services to a multiplicity of customer terminals via a common tariff distribution mechanism;

and selectively varying a respective tariff depending on an operational condition of a respective service.

5

15. A method according to any one of the preceding claims, including communicating different tariffs having different respective volatilities to different respective ones of the multiplicity of customer terminals.

10 16. A method of operating a communications network, including

calculating for each of a multiplicity of customers, using a selected one of a plurality of different tariffs, charges for the use of network resources by a respective customer terminal attached to the network,

measuring the loading of network resources, and

15 varying one or more of the plurality of different tariffs in dependence upon the loading of the network resources, and in which different ones of the plurality of different tariffs have different respective volatilities.

17. A method of operating a communications network in which at a point of
20 access to the network a single blocking test only is applied to traffic entering the network .

18. A method of operating a communications network comprising:

a) communicating tariff data to a user terminal connected to the network;

25 b) calculating at the user terminal using the tariff data a charge for traffic communicated between the network and the terminal and making a payment;

c) sampling part only of the traffic communicated between users and the network and for the sampled traffic comparing any payments made by users and the payment due according to the tariff .

30

19. A method of operating a communications network comprising;

a) establishing contracts between network users and a network operator and storing user contract data;

- b) sampling part only of the traffic to or from a user on the network;
- c) comparing sampled traffic with traffic contracted for by the user; and
- d) amending the user status when a discrepancy between the sampled parameters and the contracted parameters is detected.

5

20. A method according to claim 19, in which the step of establishing contracts between network users and the network operator includes making an advance payment for network usage.

10 21. A method according to claim 19 or 20, in which the step of amending the user status includes fining the user.

22. A method according to claim 19, in which in step (a) the user transfers a deposit to the network operator, which deposit is debited in step (d) when the
15 discrepancy between the sampled parameters and the contracted parameters is detected.

23. A method according to any one of the preceding claims, in which the communications network is a network supporting a packet-based internetworking
20 protocol.

24. A communications network arranged to operate by a method according to any one of the preceding claims.

25 25. A customer terminal adapted for use in a method according to any one of the preceding claims.

26. A customer terminal for use in a communications network, the customer terminal including;

30 a network interface which in use receives tariff information via a communications network;

a store programmed with tariff information received at the said interface;

a meter for measuring use by the customer terminal of the network to which the tariff applies; and

a processor connected to the said meter and to the store and arranged to calculate using the said tariff information a network usage charge.

27. A method of operating a communications network substantially as described
5 with respect to the accompanying drawings and in the accompanying paper.

28. A communications network substantially as described with respect to the accompanying drawings and in the accompanying paper.

10

29. A method of operating a communications network comprising

a) at a customer terminal measuring network usage;

b) communicating network usage data from the customer terminal to the network operator; and

15

c) the network operator sampling part only of the traffic communicated between a customer terminal and the network and for the sampled traffic comparing the network usage with the network usage data from the customer terminal and thereby detecting any discrepancy.

20 30. A method according to any one of claims 1 to 10 including communicating encrypted tariff data to the customer terminal, and decrypting the said tariff data at the customer terminal.

31. A method of operating a communications network including;

distributing a tariff via the communications network to a multiplicity of
25 customer terminals connected to the communications network,

measuring at a customer terminal use by the customer terminal of network resources; and

calculating, using the results of the said step of measuring together with the said tariff, a charge for use by the customer terminal of the network to which
30 the tariff applies.

32. A method of operating a communications network, including automatically varying, depending on network loading as detected at a customer terminal, a tariff for network usage by a customer terminal.

THIS PAGE BLANK (USPTO)




Figure 1

THIS PAGE BLANK (USPTO)
